Probabilistic Reachability for Stochastic Hybrid Systems: Theory, Computations, and Applications

Alessandro Abate

Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2007-132
http://www.eecs.berkeley.edu/Pubs/TechRpts/2007/EECS-2007-132.html

November 7, 2007

Copyright © 2007, by the author(s).
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

Probabilistic Reachability for Stochastic Hybrid Systems: Theory, Computations, and Applications

by

Alessandro Abate

Laurea (Università degli Studi di Padova) 2002
M.S. (University of California, Berkeley) 2004

A dissertation submitted in partial satisfaction
of the requirements for the degree of

Doctor of Philosophy

in

Electrical Engineering and Computer Sciences

in the

GRADUATE DIVISION

of the

UNIVERSITY OF CALIFORNIA, BERKELEY

Committee in charge:

Professor Shankar S. Sastry, Chair
Professor Pravin Varaiya
Professor Claire J. Tomlin
Professor Steven N. Evans

Fall 2007

The dissertation of Alessandro Abate is approved:

Professor Shankar S. Sastry, Chair	Date
Professor Pravin Varaiya	Date
Professor Claire J. Tomlin	Date
Professor Steven N. Evans	Date

University of California, Berkeley

Fall 2007

Probabilistic Reachability for Stochastic Hybrid Systems:
Theory, Computations, and Applications

Copyright © 2007

by

Alessandro Abate

Abstract

Probabilistic Reachability for Stochastic Hybrid Systems:
Theory, Computations, and Applications

by

Alessandro Abate
Doctor of Philosophy in Electrical Engineering and Computer Sciences
University of California, Berkeley
Professor Shankar S. Sastry, Chair

Stochastic Hybrid Systems are probabilistic models suitable at describing the dynamics of variables presenting interleaved and interacting continuous and discrete components.

Engineering systems like communication networks or automotive and air traffic control systems, financial and industrial processes like market and manufacturing models, and natural systems like biological and ecological environments exhibit compound behaviors arising from the compositions and interactions between their heterogeneous components. Hybrid Systems are mathematical models that are by definition suitable to describe such complex systems. The effect of the uncertainty upon the involved discrete and continuous dynamics—both endogenously and exogenously to the system—is virtually unquestionable for biological systems and often inevitable for engineering systems, and naturally leads to the employment of stochastic hybrid models.

The first part of this dissertation introduces gradually the modeling framework and focuses on some of its features. In particular, two sequential approximation procedures are introduced, which translate a general stochastic hybrid framework into a new probabilistic model. Their convergence properties are sketched. It is argued that the obtained model is more predisposed to analysis and computations.

The kernel of the thesis concentrates on understanding the theoretical and computational issues associated with an original notion of probabilistic reachability for

controlled stochastic hybrid systems. The formal approach is based on formulating reachability analysis as a stochastic optimal control problem, which is solved via dynamic programming. A number of related and significant control problems, such as that of probabilistic safety, are reinterpreted with this approach. The technique is also computationally tested on a benchmark case study throughout the whole work. Moreover, a methodological application of the concept in the area of Systems Biology is presented: a model for the production of antibiotic as a component of the stress response network for the bacterium *Bacillus subtilis* is described. The model allows one to reinterpret the survival analysis for the single bacterial cell as a probabilistic safety specification problem, which is then studied by the aforementioned technique.

In conclusion, this dissertation aims at introducing a novel concept of probabilistic reachability that is both formally rigorous, computationally analyzable and of applicative interest. Furthermore, by the introduction of convergent approximation procedures, the thesis relates and positively compares the presented approach with other techniques in the literature.

Professor Shankar S. Sastry, Chair Date

Acknowledgements

This dissertation is the final outcome of the work of a single person, but benefits from the contributions of many. I would like to acknowledge the help I have had in the past few years of graduate school at Berkeley.
Shankar Sastry, my advisor, has granted me the unconditional intellectual freedom to pursue my ideas and intuitions, the enthusiasm to aim higher, and the financial tranquillity that enabled me to focus on my work.
Maria Prandini and John Lygeros have contributed to this project in first person, dispensing their attention and their knowledge of Hybrid Systems on me. I greatly benefited from John's expertize and Maria's technical rigor.
Claire Tomlin has always been ready to give me extremely effective feedback on my work, and I am excited about our upcoming collaboration.
Steven Evans, in the department of Statistics, and Pravin Varaiya in EECS have always been ready to help me and gave me competent and thorough feedback as members of my committee.
I have learned a lot from my collaboration with the CS Lab at SRI International, in particular with Ashish Tiwari.
Laurent El Ghaoui and Slobodan Simic have been reference figures for me, as well as collaborators, especially during my Master years.
Indirectly, I have had the chance to benefit from the teaching and the exchange of ideas with exceptional people at Berkeley: Alberto Sangiovanni-Vincentelli in EECS, David Aldous and Mike Jordan in Statistics.
Saurabh Amin has proven to be a great collaborator on this research project, and part of the results presented here are his own. Alessandro D'Innocenzo, Aaron Ames, Minghua Chen and Ling Shi have shared many hours of discovery and exciting research with me. I would like to acknowledge the contribution of the control people in Cory 333, especially my fellows Luca Schenato, Bruno Sinopoli and John Koo.
On a personal note, I would like to gratefully thank the people at Berkeley that I have had the privilege to spend my best time with, and which have greatly influenced my own growth during graduate school. In particular, a heartful thank you to the people of the IISA and, foremost, to Alvise Bonivento, Fabrizio Bisetti, Alberto Diminin, and Alessandro Pinto.
The concluding acknowledgment, without adding any obvious and granted explanations or justifications for it, goes to my family and to Lynn. I know that they understand what I feel when I write these words: grazie.

Alessandro Abate
Berkeley, CA, Summer/Fall 2007

Contents

List of Figures

Introduction

Hybrid Systems (HS) are dynamical models with interacting continuous and discrete components. They naturally model continuous systems with phased, multi-modal operation, or with fault recovery procedures; hierarchical, logic-based, quantized-control systems; embedded, software and networked control of real-time, physical systems; and systems with heterogeneous models of computations. Seminal work in this research arena can be found as early as in [Witsenhausen, 1966].

This modeling framework has attracted a lot of attention in the Systems Theory and Control Engineering community because of its twofold value: first, the theoretical interest that such models promote; second, the wealth of successful practical arenas and studies resulting from their applications [Mariton, 1990; Varaiya, 1993; Brockett, 1993; Ramadge and Wonham, 1989].

On the one hand, it is in fact of great interest to systematically build up a mathematical framework for the analysis and the control of compound continuous and discrete systems: the understanding of the unique structural characteristics and the general behaviors associated with these models has allowed the exploitation of their full modeling capabilities. Also, the literature has witnessed an unprecedented confluence of a system/analytical approach to these models [Lygeros, 1996; Branicky, 1995], and a computer-science/computational perspective on them [Puri, 1995; Alur *et al.*, 1993; Alur and Dill, 1994; Henzinger *et al.*, 1998]: this has provided original attention to aspects at the intersection of the two areas. This has further spurred the study, the analysis and the verification of descriptive, implementable, computationally efficient and practically applicable models and control synthesis techniques.

On the other hand, also thanks to the advances of the theoretical understanding of these models, hybrid system models have found exciting application arenas

in transportation [Varaiya, 1993; Lygeros and Godbole, 1997; Lygeros *et al.*, 1998], automotive and air traffic control [Tomlin *et al.*, 1998b; Glover and Lygeros, 2004; Hu *et al.*, 2003; Prandini *et al.*, 2000], robotics (mechanical systems undergoing impacts or tracking of maneuvering targets, for instance), hierarchical systems [Lygeros, 1996] for safe and optimal control synthesis, industrial processes (like manufacturing, resource allocation, or fault-detection models from operations research) [Mariton, 1990], financial applications [Davis, 1984; Davis, 1993; Glasserman and Merener, 2003], networks (power, and telecommunication ones) [Hespanha *et al.*, 2001; Hespanha, 2004; Abate *et al.*, 2006a], biological and ecological systems [Lincoln and Tiwari, 2004; Alur *et al.*, 2001].

The first part of the dissertation (chapter 1, section 1.1) introduces the notion of deterministic Hybrid System, adhering to what is accepted as one of more general frameworks, that of the *hybrid automaton*. A number of structural and dynamical properties of interest in the proceeding of the thesis are described.

✠

Stochastic Hybrid Systems (SHS) have attracted the interest of the System Theory community only relatively recently. However, the results for the models that are currently investigated build on the shoulders of older mathematics (probability theory, theory of Markov chains, theory of uncertain system). Chapter 1 describes some theoretical investigations on this class of probabilistic models. In particular, the attention is focused on the problem of *finding a SHS model that is both general and descriptive, as well as prone to be analyzed and computed*. To start with, a thorough literature review of SHS is contained in section 1.2. Then, a very general class of Stochastic Hybrid Models (GSHS) are introduced in section 1.2.1—if not the most general, it is truly the one which encompasses all of the characteristics of HS and directly extends them to the probabilistic case. In section 1.2.2, some necessary technical concepts are defined, such as that of extended generator of a solution process of a stochastic hybrid model. Section 1.3.1 describes a methodology by which a GSHS model is approximated by a similar probabilistic model that has no spatial guards. The *events*, once due to the presence of spatial conditions, are now randomly determined by properly defined arrival processes, whose parameters are state-dependent.

The procedure can be likewise applied to subclasses of the GSHS, all the way down to purely deterministic HS. The intuitive fact that the new systems are "easier" than the old ones, albeit at the expense of introducing new random quantities, is motivated by a number of instances in section 1.3.3. In section 1.3.2, the weak convergence of the solution of the new model to a solution process of the original one is sketched under proper assumptions. Moreover, at the end of section 1.3.3, it is claimed that the obtained class of SHS models, while being quite general, is also apt to be analyzed and simulated. Section 1.4 is dedicated to the discretization in time of the above model. More precisely, by starting from the "approximated" SHS model obtained in section 1.3.1, a formulation of its dynamics by the use of random measures is proposed (section 1.4.1). This new form allows the application of a time sampling procedure, according to known integration methods, such as the first-order Euler scheme (section 1.4.2). The weak convergence of the obtained discrete-time processes to the original continuous-time ones is sketched in section 1.4.3.

✠

Reachability is an important and well investigated topic in classical control theory. The overture of chapter 2 gives an informal introduction to the concept, a qualitative interpretation for the stochastic models under study and a perspective of the (recently performed) work in the literature (section 2.1).

The prime objective of the dissertation is to introduce a new such concept for a general class of Stochastic Hybrid Systems. The model under the study is the one obtained at the end of chapter 1, with the addition of a control structure. This model is recalled and reframed in section 2.1.1.

The theoretical part of this section 2.2 unfolds as follows. The notion is introduced in section 2.2.1, where two alternative interpretations of the concept are suggested. The formal approach is based on formulating reachability analysis as a stochastic optimal control problem in section 2.2.2, which is solved via dynamic programming (DP) in section 2.2.3. A number of related and significant control problems, such as that of probabilistic safety or that of regulation (which subsumes the steady-state analysis in section 2.2.4), are reinterpreted via this approach and discussed in section 2.2.5, 2.2.6 and 2.2.7. The practical algorithmic solution of the dynamic programming

scheme assumes a discrete framework, both in time and in space: in section 2.2.8.1, a state space discretization procedure is introduced, its properties illustrated, and its convergence properties proved in section 2.2.8.2.

The technique is computationally tested in section 2.3 on a benchmark case study proposed in past literature: that of temperature regulation in a number of rooms, by a collection of thermostat-controlled heaters (section 2.3.1). The control synthesis problem is solved in section 2.3.2 for the single-room and the multiple-rooms case. The aforementioned concept of regulation is tested on this benchmark in section 2.3.2.2, and the convergence of the state space gridding shown in section 2.3.3. Also, the use of a number of techniques are proposed, to partly mitigate the curse of dimensionality that affects the solution of a DP, and which is common with the other approaches in literature. This is achieved in sections 2.3.4 and 2.3.4.1.

Finally, a methodological application of the concept in the area of Systems Biology is presented in section 2.4.1: a model for the production of antibiotic as a component of the stress response network for the bacterium *Bacillus subtilis* is described in section 2.4.2. The model allows one to reinterpret the survival analysis for the single bacterial cell as a probabilistic safety specification problem (see section 2.4.3), which is then studied and computed by the aforementioned technique in section 2.4.4.

This part ends with the exploration of possible future research avenues, in section 2.5.

The dissertation places particular emphasis on the connection of the present effort with other related work in the literature, both at a foundational level (chapter 1), and at a more technical level (section 2.2.9 in chapter 2). The collection of a thorough network of exhaustive references (see Bibliography) should help the reader with the comparisons and give her or him more pointers to the foundational and the adjacent work.

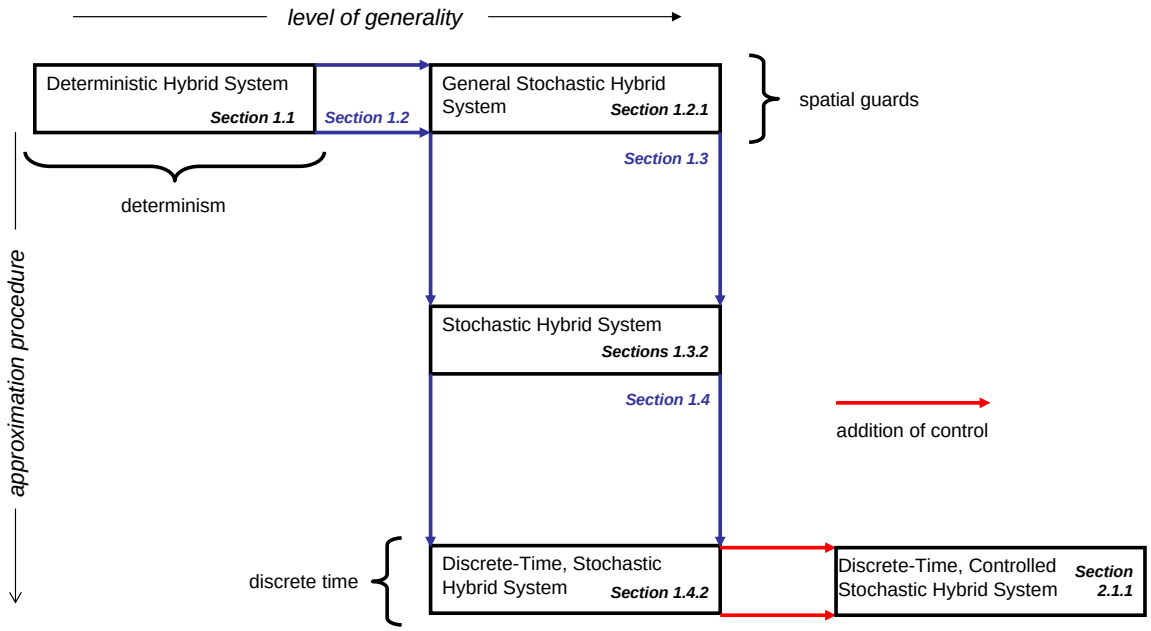

Figure 1: Dependency chart for Sections and corresponding Hybrid System Models.

Chapter 1

Modeling

1.1 Deterministic Hybrid System Model

A Hybrid System is described by a formal mathematical model. Such a model can be defined in a number of different ways, each highlighting particular features of the system under study or focusing on particular points of the structure or being more or less synthetic. In particular, some models (possibly coming from Computer Science) stress the "discrete" features of the system. Others instead, mainly relating to Systems Theory, focus on the "continuous" dynamics and behaviors.

In this dissertation the template dynamical model of choice is the *hybrid automaton* [Lygeros *et al.*, 2003; Lygeros, 2004a]. This choice comes, for the sake of homogeneity and uniformity, from the ease to extend such a mathematical entity to the controlled and the probabilistic cases. Let us reiterate that other models are similar and as meaningful as the one presented. The author, for instance, has worked on a related framework [Abate *et al.*, 2005; Abate *et al.*, 2006a].

Definition 1 (Deterministic Hybrid System)**.** *A* Deterministic Hybrid System *(HS) is a collection* $\mathscr{H} = (\mathcal{Q}, E, D, \Gamma, A, R)$*, where*

- $\mathcal{Q} = \{q_1, q_2, \ldots, q_m\}$ *is a finite set of discrete modes;*
- $E = \{e_{i,j}, i, j \in \mathcal{Q}\} \subseteq \mathcal{Q} \times \mathcal{Q}$ *is a set of edges, each of which is indexed by a pair of modes; given an edge* $e_{i,j}, i = \mathfrak{s}(e)$ *is its* source *and* $j = \mathfrak{t}(e)$ *its* target*;*

- *$D = \{D_1, D_2, \ldots, D_m\}$ is a set of domains, each of which is associated with a mode. Let us assume that $D_q \subseteq \mathbb{R}^n, n < \infty, \forall q \in \mathcal{Q}$. The hybrid state space is introduced as $\mathcal{S} = \bigcup_{q \in \mathcal{Q}} q \times D_q$ (see Remark 1);*

- *$A = \{a_q, q \in \mathcal{Q}\}, a_q : \mathcal{Q} \times D \to D$, is the set of vector fields, which are assumed to be Lipschitz. Each vector field characterizes the continuous dynamics in the corresponding domain, which evolve in continuous time;*

- *$\Gamma = \{\gamma_{i,j} \subset D_i\} \subset \mathcal{S}, j \neq i \in \mathcal{Q}$, is the guards set, a subset of the state space. They represent boundary conditions and are associated with an edge: $\forall i, j \in \mathcal{Q} : \gamma_{i,j} \in \Gamma, \exists\, e_{i,j} \in E$;*

- *$R : \mathcal{Q} \times \mathcal{Q} \times D \to D$ is a reset function, associated with each element in Γ (or, equivalently, to each edge): with the point $s = (i, x) \in \gamma_{i,j}$ is associated a reset function $R(j, (i, x))$).*

The initial condition for the hybrid solution process of the above model, will be taken from a set of hybrid values $\mathtt{Init} \subseteq \mathcal{S}$. □

Remark 1 (On the Structure of the HS)**.**

- *The pair $(\mathcal{Q}, E)$ characterizes the discrete structure of the hybrid system, that of a finite-state automaton.*

- *The* Hybrid State Space *$\mathcal{S}$ is defined as the disjoint union of the domains pertaining to each mode, $\mathcal{S} = \bigcup_{q \in \mathcal{Q}} \{q\} \times D_q$. A point in the state space will be a pair $s = (q, x)$, where $q \in \mathcal{Q}$ and $x \in D_q$. Similarly, as formally described in Algorithm 1, a hybrid trajectory will be made up of two components, a discrete one and a continuous one, each dwelling in its respective subspace, the first a discrete set of modes, the second a set of continuous domains, subsets of Euclidean spaces.*

- *Let us assume the guard set is* forcing*, that is, once it gets hit, it instantaneously elicits a jump. This affects the semantics of the model (see Algorithm*

1). More generally, guards may just enable a jump [Lygeros, 2004a]. This introduces issues of non-determinism, the details of which are not covered here;[1] *incidentally, this issue can be properly handled by stochastic models—see and page 19 and 1.3.3.* □

We have used freely the word "process," or "trajectory" of $\mathscr{H}$. The following introduces a formal definition of these concepts.

Definition 2 (Hybrid Time Set)**.** *A hybrid time set $\tau = \{I_k\}_{k\geq 0}$ is a finite or infinite sequence of intervals $I_k = [t_k, t'_k] \subseteq \mathbb{R}$ such that*

1. *I_k is closed if τ is infinite; I_k might be right-open if it is the last interval of a finite sequence τ;*

2. *$t_k \leq t'_k$ for $k > 0$ and $t'_{k-1} = t_k$ for $k > 1$.* □

The length $t'_k - t_k$ of every interval I_k denotes the dwelling time in a discrete location of the hybrid flow, while the extrema t_k, t'_k specify the switching instants. Let us stress that the above set is ordered; hence, it makes sense to use notations such as $t_k \leq t'_k$, as we shall do throughout the work.

In the succeeding work, boldface shall be used to denote trajectories or executions, and normal typeset to denote sample values. This convention will also be used for the processes, solutions of the stochastic models. Also, often the starting time for the solution of a (S)HS will be taken for simplicity to be $t_0 = 0$.

A *hybrid trajectory*, or *hybrid flow*, is a pair $(\mathbf{s}, \tau)$, where the first component is the hybrid state $\mathbf{s} = (\mathbf{q}, \mathbf{x}) : \tau \to \mathcal{S}$, that describes the evolution of the continuous part $\mathbf{x}$ and the discrete part $\mathbf{q}$ by means of (possibly multi-valued) functions defined on the hybrid time set τ and having value on $\mathcal{S}$.

Finally, an *hybrid execution* is a pair $(\mathbf{s}, \tau)$ which can be algorithmically described as follows:

Definition 3 (Hybrid Execution)**.** *Consider an HS $\mathscr{H} = (\mathcal{Q}, E, D, \Gamma, A, R)$. A trajectory pair $\mathbf{s}(t) = (\mathbf{q}(t), \mathbf{x}(t))$ with values in $\mathcal{S}$ is an execution of $\mathscr{H}$ associated with*

[1] For further details, such as sufficient conditions to prevent non-determinism, refer to [Abate *et al.*, 2006b].

an initial condition $(\mathbf{q}(t_0), \mathbf{x}(t_0)) \in \texttt{Init}$ *if it is obtained according to the following scheme:*

Algorithm 1.

1. *At starting time* $t_0 \geq 0$, *pick* $(\mathbf{q}(t_0), \mathbf{x}(t_0)) \in \texttt{Init}$, *set* $k = 0$, $\tau = \varnothing$;
2. *Extract a continuous trajectory* $\mathbf{x}(t)$ *from the vector field and with initial condition* $x(t_k)$ *until possibly a guard is hit: namely until time* $t'_k \in [t_k, \infty)$ *such that* $\mathbf{x}(t'_k) \in \gamma_e$, *where* $\mathfrak{s}(e) = q(t_k)$;
3. *If* $t'_k = \infty$,

   add $I_k = [t_k, \infty)$ *to* τ *and exit the algorithm;*
4. *Else* *add* $I_k = [t_k, t'_k]$ *to* τ;

   define $\mathfrak{t}(e) = \mathbf{q}(t_{k+1})$ *and* $\mathbf{x}(t_{k+1}) = R(\mathfrak{t}(e), \mathfrak{s}(e), x(t'_k))$;

   increment k *and go to line 2.* □

In what follows, let us define an *event* to be a discrete state transition associated with the hitting of the guard set by the hybrid trajectory.

Remark 2 (Blocking Conditions)**.** *In Algorithm 1, it has been implicitly assumed that either an event happens in finite time (whenever the hybrid execution intersects the guard set), or that the execution dwells indefinitely inside a domain. This excludes the case when a trajectory exits a domain without necessarily hitting a guard, which in the parlance is known as a* blocking condition*. It is possible to exclude this from happening if it is assumed that the boundary of each domain is included in the corresponding guard set, and noticing that the image of the reset map is a subset of the domain set. For more details on blocking, explicit conditions to prevent it, and other structural issues for deterministic HS, refer to [Lygeros et al., 2003; Abate et al., 2006b].* □

Hybrid Systems pose a problem which is unknown in the simpler setting of dynamical systems, that of Zeno dynamics. In simple terms, Zeno behaviors happen when, in a bounded time interval, the hybrid trajectory jumps between specific domains infinitely many times. More precisely, consider the following definition:

Definition 4 (Zeno Behavior)**.** *A hybrid system $\mathscr{H}$ is Zeno if for some execution $\mathbf{s}(t), t \in \tau : \tau \to \mathcal{S}$ of $\mathscr{H}$ there exists a finite constant t_∞ (called the Zeno time) such that*

$$\lim_{i\to\infty} t_i = \sum_{i=0}^{\infty}(t_{i+1} - t_i) = \sum_{i=0}^{\infty} I_i = t_\infty.$$

the execution $\mathbf{s}(t), t \in \tau$ is called a Zeno execution. □

Classification of Zeno Behavior. The definition of a Zeno execution results in two qualitatively different types of Zeno behavior. They are defined as follows: for an execution $\mathbf{s}(t), t \in \tau$ that is Zeno, $\mathbf{s}(t), t \in \tau$ is

Chattering Zeno: If there exists a finite constant C such that $t_{i+1} - t_i = 0$ for all $i \geq C$.

Genuinely Zeno: If $\forall i \in \mathbb{N}, \exists k > 0 : t_{i+k+1} - t_{i+k} > 0$.

The difference between these two classes is especially prevalent in their detection and elimination. Chattering Zeno can be studied by considering solutions *a-la-Filippov* [Sastry, 1999; Khalil, 2001], that is looking at the specific part of the HS that is associated with chattering behaviors as originating from a vector field with discontinuous righ-hand side. Chattering Zeno is also relatively easy to detect [Zhang *et al.*, 2001; Ames and Sastry, 2005b]. Genuinely Zeno executions instead are much more complicated in their behavior, as well as in their detection. It has been only recently that sufficient conditions have been developed to prove their existence [Heymann *et al.*, 2002; Ames *et al.*, 2005; Heymann *et al.*, 2005; Ames *et al.*, 2007]. More precisely, [Ames *et al.*, 2005], introduces a simpler but effectively equivalent definition of Genuine Zeno ("$\forall i \in \mathbb{N}, t_{i+1} - t_i > 0$"), and develops sufficient conditions for a simple class of HS, namely *diagonal, first quadrant hybrid systems*. The results are based on a study of the hybrid dynamics in a neighborhood of the Zeno point (or Zeno *equilibrium*, defined as a hybrid point on which Zeno behavior happens), and on the use of a Poincaré-like map, which allows the study of the behaviors in discrete time. The work in [Ames *et al.*, 2007] extends the previous results to the nonlinear case by the use of a *flow box theorem*-like argument. Furthermore, the notion of *stable Zeno point* is introduced, and Zeno behaviors are related to exponentially stable Zeno equilibria.

Let us now introduce a well known modeling example. With the exception of the last section 2.4.1, where a case study from the field of Systems Biology will be introduced, we shall adhere to the following application instance throughout the dissertation by first further developing the model, and then by using it on a computational case study. This use is motivated by some benchmarks, which recently appeared in the literature [Fehnker and Ivančić, 2004], and which will be targeted for computational test comparisons.

Example 1 (Thermostat)**.** *In order to model the temperature dynamics of a room with a heater controlled by a thermostat, let us introduce the hybrid model* $\mathscr{H}^{th} = (\mathcal{Q}, E, D, \Gamma, A, R)$*, with*

- $\mathcal{Q} = \{ON, OFF\}$;
- $E = \{(ON, OFF), (OFF, ON)\}$;
- $D = \{D_{ON} = \{x \in \mathbb{R} : x \leq 80\}, D_{OFF} = \{x \in \mathbb{R} : x \geq 70\}\}$;
- $\Gamma = \{G_{ON} = \{x \in \mathbb{R} : x \geq 79\} \cap D_{ON}, G_{OFF} = \{x \in \mathbb{R} : x \leq 71\} \cap D_{OFF}\}$;
- $A = \{a_{ON}(x) = -\alpha(x - 90), a_{OFF}(x) = -\alpha x\}$, *where* $\alpha > 0$;
- $R = \{\mathit{id}(x)\}$.

The set of initial conditions is taken to be the complement of the guard set, with respect to the domain set, $\mathtt{Init} = D \backslash \Gamma$*. It is assumed that the evolution is in continuous time. In figure 1.1 (right) the hybrid automaton model for the system is represented. In figure 1.1 (left), a simple simulation is implemented, where the time horizon is 10 seconds,* $\alpha = 0.1$*, the initial condition* $\mathbf{s}(0) = 75^o F$*. Recall that the guards have the semantics of "forcing" transitions.* □

Remark 3 (Switching and Hybrid Systems)**.** *The reader should ponder over the difference between the framework of "switching" systems,*[2] *and that of "hybrid" systems.*

[2] To be further precise, the literature distinguishes between autonomous "switching" systems and controlled "switched" systems. However, we do not further pursue this difference in the remaining of this dissertation.

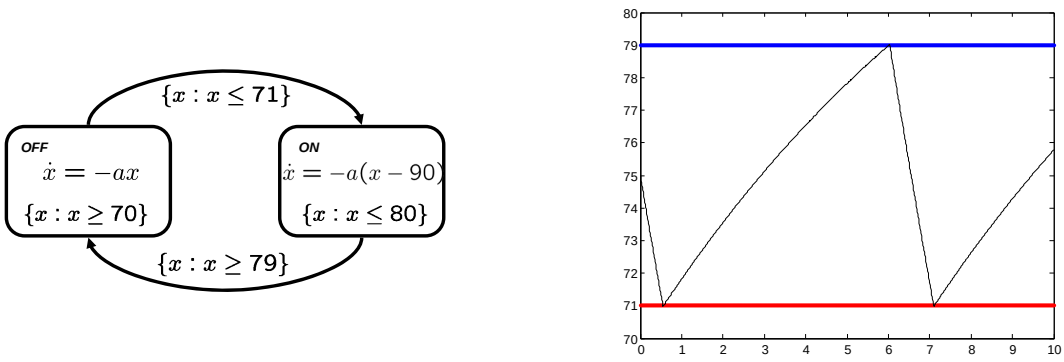

Figure 1.1: Left: hybrid automaton model for the deterministic thermostat system. Right: MATLAB simulation of the thermostat model (the horizontal axis represents time in seconds, the vertical temperature in Fahrenheit).

The first category is characterized by event conditions "in time," that is a priori defined through a sequence of jumping times $\{t_k\}_{k\in\mathbb{N}}$. *The second modeling framework, instead, specifies possible event conditions in terms of the variables of the model, by introducing a guard set. The event times are then specified on the single trajectory, and hence the sequence of these times varies depending on the single initial condition.*

It is intuitive that hybrid models are intrinsically more complicated than switched ones. Often, to prove properties of a hybrid system, it is worth considering a switched model, which is a simulation [Tanner and Pappas, 2002] of it and, as such, contains all of its behaviors. Properties are then proved on the simulation and translated back to the original hybrid model [Abate and Tiwari, 2006]. □

A note on Controlled Deterministic Hybrid System Model. It is legitimate to introduce the presence of a control structure on the HS model defined above. However, given our interest in developing control problems only on a discrete-time setting, we defer such extension to chapter 2. The interested reader is referred to [Abate *et al.*, 2006b] for a definition and a discussion of related modeling issues.

1.2 Probabilistic Generalizations

Motivations. The introduction of a stochastic analog to the model presented in section 1.1 is mainly motivated by two arguments:

To begin with, a probabilistic model is more general, in the behaviors for which it allows, than its deterministic counterpart. Indeed, a deterministic model can be thought of as being a possible implementation/instance of a probabilistic one. That is, the deterministic behaviors are "contained" in the set of stochastic ones [Polderman and Willems, 1998]. The Systems Theory literature has witnessed a progressive generalization in the analysis of dynamical systems. From deterministic models, to uncertain models [Bertsekas, 1971] (where parameters or signals are known "within bounds"), the study reaches a general breadth when models that are explicitly probabilistic are analyzed and understood. This often comes with the burden of more and deeper technicalities that are required in the analysis. In a number of instances, though, stochastic models may appear to be more manageable than deterministic ones, as it we further discuss in section 1.3.3.

Furthermore, the introduction of explicitly probabilistic models has been considered necessary in a number of applicative instances, where the uncertainty entering the system cannot be simply "averaged out," when the knowledge of the system is just too coarse, or when it is evident/manifest that some stochastic mechanisms play a role in the system under study. Known instances of these last arguments are models drawn from finance [Cont and Tankov, 2004], models describing air traffic control applications [Prandini *et al.*, 2000; Blom and Lygeros, 2006], or models describing specific biological phenomena [Gillespie, 1976; Gillespie, 1977] (especially when only a limited number of entities are involved).

Levels of Generalization. The deterministic model $\mathscr{H} = (\mathcal{Q}, E, D, \Gamma, A, R)$ is made up by elements that could be embedded with some randomness. More precisely, probabilistic terms can be introduced in $\mathscr{H}$ at the following levels:

1. $(\mathcal{Q}, E)$, the underlying discrete structure: rather than having a state-automaton structure, let us think of having a continuous-time Markov-chain sort of relationship between the modes. The jumps may be due to some transition intensities,

rather than events defined as conditions by R;

2. A, the continuous dynamics: let us introduce probabilistic continuous dynamics, in terms of a stochastic differential equation, rather than the deterministic ODE;

3. R, the discrete resets: let us have probabilistic resets, described by stochastic kernels, rather than the deterministic functions in R;

4. `Init`, the set of initial conditions: the starting state may be sampled from a probabilistic distribution, rather than being deterministically picked within the set of initial conditions.

Let us introduce a rather simple modification of Example 1, which is endowed with probabilistic continuous dynamics.

Example 2 (Stochastic Thermostat)**.** *Let the hybrid model $\mathscr{S}^{th} = (Q, E, D, \Gamma, A, \Sigma, R)$ be made up of the same elements of $\mathscr{H}^{th}$, except for the continuous dynamics, which are now enhanced with a diffusion term and thus described by a stochastic differential equation (SDE) [Arnold, 1992; Oksendal, 1998] so that, for $q \in Q, t \geq 0$,*

$$d\mathbf{x}(t) = a_q(\mathbf{x}(t))dt + b_q^T b_q d\mathbf{w}(t),$$

where $\mathbf{w}(t)$ is a one-dimensional standard Wiener process. The following has been introduced:

- $\Sigma = \{\sigma_{ON}(x) = \sigma_{OFF}(x) = \sqrt{8}\}$.

In figure 1.2 (right) the hybrid automaton model for the system is shown. In figure 1.2 (left), a simple simulation is implemented, where the time horizon is 10 seconds, $\alpha = 0.1$, the initial condition $s(0) = 75^o F$. Recall again that the guards are assumed to be "forcing" an event. □

Literature Review on Stochastic Hybrid Systems. The framework of Stochastic Hybrid System (SHS) is rather general and encompasses other mathematical models that have been widely investigated in the literature. It is often inspiring to study, in generality, how properties that are well understood in the context of SDE's or

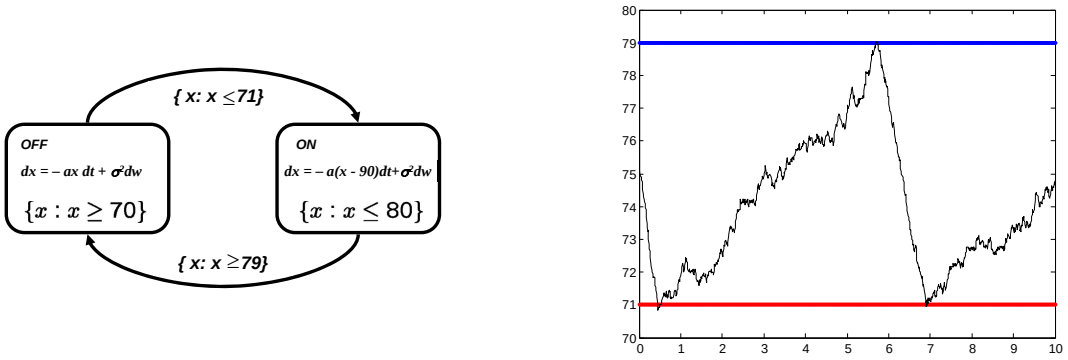

Figure 1.2: Left: hybrid automaton model for the stochastic thermostat system. Right: MATLAB simulation of the thermostat model (the horizontal axis represents time in seconds, the vertical temperature in Fahrenheit).

Lévy processes (first hitting times, occupation measures, martingale properties, for instance) can be exported to a "switching" or hybrid case.

A body of literature has focused on systems with Markovian switchings (see remark 3), i.e. models that progress deterministically in the continuous dynamics, while jumping according to some Poisson arrivals and randomly switching according to an underlying Markov chain structure [Mariton, 1990]. A wealth of research has been spent on proving stability properties for systems of this sort. Some work has focused on weak concepts of stability [Geromel and Colaneri, 2006; Yuan and Lygeros, 2005b], other on stronger notions [Yuan and Lygeros, 2005a; Bolzern *et al.*, 2006b; Bolzern *et al.*, 2006a]. Often the approach has been that of extending results or techniques developed for the deterministic case [Branicky, 1994; Branicky, 1995; Kourzhanski and Varaiya, 1996; Liberzon *et al.*, 1999; Liberzon, 2003], such as the use of a common Lyapunov function [Chatterjee and Liberzon, 2004], or that of multiple Lyapunov functions, with switching conditions. These conditions are often interpreted as averaging criteria for the probabilistic case [Abate *et al.*, 2004; Chatterjee and Liberzon, 2006], or translated into supermartingale [Durrett, 2004; Billingsley, 1995] conditions on the switching processes.

A seminal work which has heavily influenced the literature on SHS is that by [Davis, 1993]. This work, introducing the *Piecewise Deterministic Markov Processes* (PDMP) framework, has pinned down a number of technical issues for these models,

such as their Markovian nature, their topological structure, the rigorous definition of the associated extended generator, as well as some issue in their optimal control (both in the continuous dynamics, as well as in their switchings). The models are deterministic in their continuous evolution, can switch either because of a (deterministic) spatial condition, or because of some transition intensities. Furthermore, they are reset randomly upon jumping between different modes of operation. See also [Davis, 1984] for further details.

The rigorous work in [Ghosh *et al.*, 1992; Ghosh *et al.*, 1997] has investigated the problem of optimal control for the case of *switching diffusions*. This model describes the evolution of a process depending on a set of stochastic differential equations, among which the process jumps according to some state-dependent transition intensities. Notice that, on the one hand, the resets are identical, while on the other the introduced controls are randomized.

One of the first efforts to introduce a formal model explicitly for SHS was attempted in [Hu *et al.*, 2000], where a system evolving according to probabilistic dynamics possibly jumps between different operating modes according to some (deterministic) conditions on the state space, and resets probabilistically according to some distributions.

The hybrid structure in [Hu *et al.*, 2000] has been blended with the PDMP approach [Davis, 1993] in [Bujorianu and Lygeros, 2004a; Bujorianu and Lygeros, 2004c; Bujorianu and Lygeros, 2004b; Bujorianu and Lygeros, 2006]. In this line of research a general model for SHS is introduced and its structural and dynamical properties studied. Because of its generality, this will be the model introduced in section 1.2.1 and initially worked on. Further investigations by the same authors have focused on the optimal control of these models and some reachability issues related to them.

Other very general SHS models have been investigated, along with their Markovian properties, in [Ghosh and Bagchi, 2004; Blom, 2003], and their simulations has been studied in [Blom and Bloem, 2004].

[Lygeros *et al.*, 2006] introduces a SHS model with delays, which will further be discussed in more details in the following. In [Yuan and Lygeros, 2006] its asymptotic stability properties are studied.

A brief overview of some SHS models, with a tentative comparison between them,

is contained in [Pola *et al.*, 2003].

1.2.1 General Stochastic Hybrid System Model

Let us introduce a general Stochastic Hybrid Systems model, first worked out in [Bujorianu and Lygeros, 2004b] and refined in [Bujorianu and Lygeros, 2006]. The choice of this particular model hinges on its generality.

For the sake of clarity, a definition is introduced that depends on a particular, but general enough, choice of the continuous dynamics, which will be characterized by SDE's.

Definition 5 (General Stochastic Hybrid System)**.** *A* General Stochastic Hybrid System *(GSHS) is a collection* $\mathscr{S}_g = (\mathcal{Q}, n, A, B, W, \Lambda, \Gamma, R^{\Lambda}, R^{\Gamma})$*, where*

- $\mathcal{Q} = \{q_1, q_2, \ldots, q_m\}, m \in \mathbb{N}$ *is a countable set of discrete modes;*
- $n : \mathcal{Q} \to \mathbb{N}$ *is a map that determines the dimension of the domain associated with each mode.*[3] *For* $q \in \mathcal{Q}$*, the domain* D_q *is the Euclidean space* $\mathbb{R}^{n(q)}$*. The hybrid state space is introduced as the disjoint union of the domains:* $\mathcal{S} = \bigcup_{q\in\mathcal{Q}}\{q\} \times D_q$*;*
- $A = \{a_q, q \in \mathcal{Q}\}, a_q : D_q \to D_q$ *is the drift term in the continuous dynamics;*
- $B = \{b_q, q \in \mathcal{Q}\}, b_q : D_q \to D_q \times D_q$ *is the* $n(q)$*-dimensional diffusion term in the continuous dynamics;*
- $W = \{\mathbf{w}_q, q \in \mathcal{Q}\}, \mathbf{w}_q$ *is an* $n(q)$*-dimensional standard Wiener process;*
- $\Lambda : \mathcal{S}\times\mathcal{Q} \to \mathbb{R}^+$ *is the transition intensity function. In particular, for* $j \neq i \in \mathcal{Q}$*,* $\lambda(s = (i, x), j) = \lambda_{ij}(x)$*;*[4,5]

[3] If $\texttt{card}(\mathcal{Q}) < \infty$, then it is possible to embed each domain $D_q, q \in \mathcal{Q}$ (and its corresponding dynamics) into their union $\bigcup_{q\in\mathcal{Q}} D_q \subset \mathbb{R}^N$, where $N = \max_{q\in\mathcal{Q}} n(q)$. Notice the potential difference with the deterministic model in terms of the extension of the domains.

[4] In general, jumps and resets into the same mode could be allowed, but for the sake of clarity and notation let us rule this out at this level—the extension to that instance is straightforward and easy to work out. Notice that the actual domain of definition of Λ can be limited to $\mathcal{S} \setminus \Gamma$, which is not done here because of the approximation procedure introduced in 1.3.1.

[5] Notice that other authors, mainly following [Davis, 1993], use a global intensity function $\lambda : \mathcal{S} \to$

- $R^{\Lambda} : \mathcal{B}(\mathbb{R}^{n(\cdot)}) \times \mathcal{Q} \times \mathcal{S} \to [0,1]$, *and denoted as* $R^{\Lambda}(j, \cdot, (i,x)) = R^{\Lambda}(\cdot | j, (i,x))$, *is a reset stochastic kernel associated with jumps elicited by* Λ*;*

- $\Gamma = \{\bigcup_{j \neq i, j \in \mathcal{Q}} \gamma_{i,j} \subset D_i\} \subset \mathcal{S}$ *represents the closed guard set of the each of the domains, where* $\gamma_{i,j}$ *are closed sets as well. It could either represent the boundary of a domain (like in the deterministic case), or only a subset of it, which deterministic jump events are associated with;*

- $R^{\Gamma} : \mathcal{B}(\mathbb{R}^{n(\cdot)}) \times \mathcal{Q} \times \mathcal{S} \to [0,1]$, *denoted as* $R^{\Gamma}(j, \cdot, (i,x)) = R^{\Gamma}(\cdot | j, (i,x))$ *is a reset stochastic kernel associated with the point* $s = (i,x) \in \gamma_{i,j}$*, which describes the reset probabilities associated with the elements in* Γ*.*

The initial condition for the stochastic solution of the above model, will be sampled from an initial probability distribution $\pi : \mathcal{B}(\mathcal{S}) \to [0,1]$. □

Remark 4 (Continuous Dynamics)**.** *The continuous dynamics (which unfold during the intervals of time when* $\mathbf{q}(t)$ *is constant), depending on elements in the sets* (A, B)*, are characterized, for any* $q \in \mathcal{Q}$*, by an SDE of the form*

$$d\mathbf{x}(t) = a(\mathbf{q}(t), \mathbf{x}(t))dt + b(\mathbf{q}(t), \mathbf{x}(t))d\mathbf{w}(\mathbf{q}(t)), \tag{1.1}$$

where $\mathbf{w}(t)$ *is a standard,* $n(q)$*-dimensional Wiener process. We have denoted* $a_q(\mathbf{s}(t)) = a(\mathbf{q}(t), \mathbf{x}(t))$*,* $b_q(\mathbf{s}(t)) = b(\mathbf{q}(t), \mathbf{x}(t))$ *and* $\mathbf{w}_q(t) = \mathbf{w}(\mathbf{q}(t))$. □

In principle, the structure of the reset kernels R^{Λ} and R^{Γ} is indistinguishable, except for their actual domain of definition, which is precisely $\mathcal{Q} \times (\mathcal{S} \setminus \Gamma)$ for R^{Λ}, and $\mathcal{Q} \times \Gamma$ for R^{Γ}.

The following is a list of assumptions on the elements of $\mathscr{S}_g$, which shall be selectively invoked in order to prove properties for the model.

Assumption 1. *Introduce the following statements:*

$\mathbb{R}^{+}$ on the whole state space, leaving the determination of the target mode to the reset function, $R^{\lambda}(\cdot, \cdot, (i,x)) = R^{\Lambda}(\cdot, \cdot | (i,x))$.

1. *Both the sets A and B are made up of functions that are globally Lipschitz continuous within each domain:*

$$\exists 0 < \mathfrak{L} < \infty : \forall q \in \mathcal{Q}, x, y \in D_q, \tag{1.2}$$

$$\|a(q,x) - a(q,y)\| + \|b(q,x) - b(q,y)\| \leq \mathfrak{L}\|x - y\|. \tag{1.3}$$

2. *Condition 1.1 , relaxed to local Lipschitz.*

3. *Both the drift and the diffusion terms are bounded in space.*

4. *The drift and the diffusion terms verify the following growth condition: there exists a positive constant $\mathfrak{C} < \infty$, such that, $\forall s = (q,x) \in \mathcal{S}$,*

$$\|a(q,x)\|^2 + \|b(q,x)\|^2 \leq \mathfrak{C}(1 + \|x\|^2).$$

5. *The initial condition, sampled from π, is independent of $w(t), t \geq 0$, of the random events coming from the distribution function (1.4), and of the probabilistic reset kernels R^Λ, R^Γ.* [6]

6. *Λ is measurable and such that, for any $q' \in \mathcal{Q}, s_0 \in \mathcal{S}$, there exists a time interval $I(q', s_0)$ which is such that the intensity function $t \to \Lambda(\mathbf{s}(t), q')$, with $\mathbf{s}(0) = s_0$ is integrable over $I(q', s_0)$.*

7. *Λ is bounded on its domain of definition.*

8. *The reset kernels R^Λ, R^Γ are Borel measurable (that is, measurable with respect to the Borel σ-algebra defined on their domain). Also, their support is bounded.*

9. *$\mathscr{S}_g$ allows no Zeno behaviors.*[7] □

[6] This condition is necessary for the existence and uniqueness of a global solution of (1.1), $\mathbf{s}(t), t \geq 0$, over any time interval [Has'minskiy, 1980; Arnold, 1992].

[7] [Davis, 1993, Assumption 24.4, Proposition 24.6] derives conditions to rule out finite "escape time" at this level, which are sufficient to exclude Zeno behaviors for $\mathscr{S}_g$. These conditions prevent any possible pathological behaviors coming from the reset kernels R^Γ and its interaction with Γ, for instance allowing only resets that are within the domains and bounded away from the guard set. It will be seen in section 1.3.3 that, given a different modeling framework for SHS, Assumption 1.7 will suffice to rule out Zeno behaviors.

The state of a GSHS is characterized by a discrete and a continuous component. The discrete state component takes on values in a countable set of modes $\mathcal{Q}$. The continuous state space in each mode $q \in \mathcal{Q}$, excluding the relative guards set, is given by a subset D_q of the Euclidean space $\mathbb{R}^{n(q)}$, whose dimension $n(q)$ is determined by the map $n : \mathcal{Q} \to \mathbb{N}$. Thus the hybrid state space is $\mathcal{S} := \bigcup_{q\in\mathcal{Q}}\{q\} \times D_q$. Let $\mathcal{B}(\mathcal{S})$ be the σ-field generated by the subsets of $\mathcal{S}$ of the form $\bigcup_q\{q\} \times A_q$, where A_q is a Borel set in D_q. It is possible to show that $\mathcal{S}$ can be endowed with a metric. A possible metric is based on the introduction of the notion of a distance. This notion is equivalent to the usual Euclidean distance metric when restricted to each domain D_q, albeit being rescaled to the unit interval. It is instead equal to one when calculated on hybrid points belonging to different discrete domains [Davis, 1993]. The topology generated by this metric shows that $(\mathcal{S}, \mathcal{B}(\mathcal{S}))$ is a Borel space , i.e. it is homeomorphic to a Borel subset of a complete separable metric space. Notice that this is not the only possible metric that can be induced on the hybrid state space.

As it shall be briefly seen (Proposition 2), the solution process for the model in Definition 5, which evolves according to the semantics of the algorithm in Definition 6, is constructed on a canonical space $\Omega = D_{\mathcal{S}}[0, \infty)$ of right-continuous $\mathcal{S}$-valued functions defined on $\mathbb{R}^+$, with left-limits (*càdlàg* functions). This space is commonly known as *Skorokhod space*. Let us not deal at this level with any possible notion of metric on this function space (the interested reader should refer to [Billingsley, 1995]). We will be working on a probability space $(\Omega, \mathscr{F}, \mathcal{P})$, made up of the sample space Ω, an associated σ-field $\mathscr{F}$ of subsets of Ω, and an induced probability measure $\mathcal{P}$ acting on the process [Davis, 1993; Durrett, 2004]. More precisely, let us endow the sample space Ω with $(\mathscr{F}_t)_{t\geq 0}$, its *natural filtration*, that is the smallest right-continuous σ-field, such that for $s \leq t, \mathscr{F}_s \subset \mathscr{F}_t$, and such that all the random variables $s(u) = (q(u), x(u))$, for a certain $u \in [0, t]$, are measurable, with respect to the induced probability on the trajectories $\mathcal{P}$ [Davis, 1993]. In particular, $\mathscr{F}_0$ includes all the $\mathcal{P}$-null states. Also, consider the sigma field $\mathscr{F} = \vee_t \mathscr{F}_t$, that is the smallest sigma field that contains all the $(\mathscr{F}_t)_{t\in\mathbb{R}^+}$. The collection $(\Omega, \mathscr{F}, (\mathscr{F}_t)_{t\in\mathbb{R}^+}, \mathcal{P})$ is called a *filtered probability space*.

A *stopping time* η on $(\Omega, \mathscr{F}, (\mathscr{F}_t)_{t\in\mathbb{R}^+}, \mathcal{P})$ is a random variable taking values in $\mathbb{R}^+ \bigcup\{+\infty\}$, such that $(\eta \leq t) \in \mathscr{F}_t, \forall t \in \mathbb{R}^+$. A process $\mathbf{M}(t), t \in \mathbb{R}^+$ is a *martingale*

of $\mathscr{F}_t$ if $E[|\mathbf{M}(t)|] < \infty$, if $\mathbf{M}(t)$ is $\mathscr{F}_t$-measurable, and if for any $s \leq t$, $E[\mathbf{M}(t)|\mathscr{F}_s] = \mathbf{M}(s)$. The concepts of super- and sub-martingale are similarly introduced. A process $\mathbf{M}(t), t \in \mathbb{R}^+$ is a *local martingale* if there is an increasing sequence of stopping times η_n, such that almost surely $\eta_n \uparrow \infty$ and the process $\tilde{\mathbf{M}}_n(t) \doteq \mathbf{M}(t \wedge \eta_n)$ is a uniformly integrable martingale, for all n.

Similar to the deterministic case (Algorithm 1), the semantics of a GSHS can be defined by the introduction of the concept of *execution*. Before defining this concept, let us introduce the notion of *survivor function*. Let $\omega^{(q,x)}(t)$ be a sample path, evolving in D_q, starting from $x = \omega^{(q,x)}(0)$. Introduce a set of functions $F_q : \mathbb{R}^+ \times \mathbb{R}^{n(q)} \to [0,1]$ for $\mathbf{s}(t) = (\mathbf{q}(t), \mathbf{x}(t)) = (q, \mathbf{x}(t))$:

$$F_q(t, \omega^{(q,x)}) = I_{\{t < t_\star(\omega^{(q,x)})\}} e^{-\int_0^t \sum_{q' \neq q, q' \in \mathcal{Q}} \lambda_{qq'}(\omega^{(q,x)}(u)) du}, \tag{1.4}$$

where, for a process starting at $(q,x) \in \mathcal{S}$, $t_\star(\omega^{(q,x)})$ is the stopping time associated with the first hitting, by the sample path $\omega^{(q,x)}$, of any guard $\gamma_{q,q'} \in \Gamma_q, q' \neq q \in \mathcal{Q}$. The first term is then related to to forced transitions. The second term describes an exponential distribution characterizing the likelihood of an event, due to the presence of the transition intensities corresponding to mode q. Notice that the quantity in the integral exists for any $t \geq 0$, by Assumption 1.6. The function $F(t, \omega^{(q,x)})$ describes the probability that no event has happened along the time horizon $[0,t]$, while the continuous motion $\omega^{(q,x)}$ unfolds within D_q, with initial condition x. Introduce the mode-dependent quantity $\widehat{\lambda}_i = \sup_{x \in D_i} \sum_{\substack{j \neq i \\ j \in \mathcal{Q}}} \lambda_{i,j}(x), i \in \mathcal{Q}$, which exists and is finite by Assumption 1.7. An execution is algorithmically defined as follows:

Definition 6 (Execution)**.** *Consider a GSHS* $\mathscr{S}_g = (\mathscr{Q}, n, A, B, W, \Lambda, \Gamma, R^\Lambda, R^\Gamma)$ *and a time horizon* $[0,T], T \in \mathbb{R}^+$*. A stochastic process* $\{\mathbf{s}(t) = (\mathbf{q}(t), \mathbf{x}(t)),\ t \in [0,T]\}$ *with values in* $\mathcal{S} = \bigcup_{q \in \mathcal{Q}} \{q\} \times D_q$ *is an execution of* $\mathscr{S}_g$*, associated with an initial condition* $s_0 = (q_0, x_0) \in \mathcal{S}$ *sampled from* $\pi \in \mathscr{P}(\mathcal{S})$*, if its sample paths are obtained according to the following algorithm:*

Algorithm 2.

set $\widehat{q} = q_0$, $\widehat{x} = x_0$, $k = 0$ *and* $T_0 = 0$*;*

while $t < T$ *do*

extract a sample path $\omega^{(\widehat{q},\widehat{x})}(t), t \in [T_k, T]$ *from the SDE in* (1.1)*, initialized on* $(\widehat{q}, \widehat{x})$*;*

extract a time $\widehat{T}$ *from the random variable* $\inf\{t > 0 | F(t, \omega^{(\widehat{q},\widehat{x})}(t)) \leq e^{-\widehat{\lambda}_{\widehat{q}} t}\}$*;*

set $s(t) = (\widehat{q}, \omega^{(q,x)}(t - T_k)), t \in [T_k, T_k + \widehat{T} \wedge T]$*;*

if $T_k + \widehat{T} < T$*, select* $(\widehat{q}, \widehat{x})$ *according to* $R(\cdot, s(T_k + \widehat{T}))$*;*

set $T_{k+1} := T_k + \widehat{T}$*;* $k := k + 1$*;*

end. □

Having introduced the concept of executions, it is now possible to seek conditions for the existence and the uniqueness of a solution of $\mathscr{S}_g$.

Proposition 1 (Existence and Uniqueness of solution of GSHS)**.** *If Assumption 1.1, 1.3, 1.5, 1.6, 1.7, 1.8 are valid, then, the GHSH in Definition 5 admits an existing and unique solution globally, that is over any time interval.*
If Assumption 1.2, 1.4, 1.5, 1.6, 1.7, 1.8 are valid, then, the GHSH in Definition 5 admits an existing and unique solution locally. □

The proof of the above statement can be directly adapted from that developed in [Ghosh and Bagchi, 2004] for a less general class of Stochastic Hybrid Systems. A uniform elliptic condition on the square of the norm of the diffusion term may allow one to relax the continuity assumption on the drift term. Furthermore, this assumption may be necessary whenever optimal control problems have to be tackled in this framework [Borkar, 1989; Ghosh and Marcus, 1995; Ghosh *et al.*, 1997; Borkar *et al.*, 1999].

The execution of a GSHS is a càdlàg function of time, i.e. continuous from the right and with left limit. It can be shown that the model in Definition 5, with unique solution evolving according to the scheme in 6 and under proper conditions, preserves the Markov property. In simple terms, the solution $\mathbf{s}(t), t \geq 0$ of the GSHS $\mathscr{S}_g$, is a Markov process if, for any $t \geq u \geq 0$ and $f \in B(\mathcal{S})$,

$$E[f(\mathbf{s}(t))|\mathscr{F}_u] = E[f(\mathbf{s}(t))|\mathbf{s}(u)].$$

This fact has relevant implications, both theoretically and computationally.

Proposition 2. *Consider the GSHS $\mathscr{S}_g$. Assume that there exists a unique solution $\mathbf{s}(t), t \geq 0$ (Proposition 1). If Assumption 1.9 holds true, then the process $\mathbf{s}(t)$ is Markov.* □

The actual proof of this statement can be found in [Bujorianu and Lygeros, 2006], where it is shown that the càdlàg and Markov properties of simpler diffusion processes can be exported to the hybrid case (including the presence of resets) by constructing *strings* of such processes, i.e. proper "concatenations" of them. Let us not pursue here further refinements of the above notions, such as that of *strong* Markov property. For more details along this line of work, please refer to [Davis, 1993] (PDMP case), [Ghosh *et al.*, 1992] (switching diffusion), [Ghosh and Bagchi, 2004] (switching diffusions with deterministic resets). From a different perspective, if the cardinality of the discrete part of the hybrid state space is finite, and each domain has the same dimension, the above model can be described by the use of *random measures*. This approach is embraced in section 1.4.1, where this modeling framework is used to prove weak convergence of time-discretization schemes. Within this approach, [Blom, 2003] shows that Markov properties are preserved under assumptions that are equivalent to those in Proposition 2.

1.2.2 Process Semigroup and Generator

Consider the space of real-valued, bounded and continuous functions f on the metrizable space $\mathcal{R}$, denoted as $C_b(\mathcal{R})$. The norm of choice is that of the sup, $\|f\| = \sup_{s\in\mathcal{R}} |f(s)|$. Consider also the space of real-valued, measurable functions on $\mathcal{R}$, $B(\mathcal{R})$.

For any $t \in \mathbb{R}^+$, define an *operator* $P_t : B(\mathcal{R}) \to B(\mathcal{R})$, for $s \in \mathcal{R}$, as

$$P_t f(s) = E_s[f(\mathbf{s}(t))].$$

Contraction and semigroup properties can be easily shown (see [Davis, 1993; Ethier and Kurtz, 1986] for more details):

$$\|P_t f\| \leq \|f\|,$$
$$P_t(P_s f)(x) = P_{t+s} f(x),\ t, s \in \mathbb{R}^+, x \in \mathcal{R}.$$

It is possible to associate to P_t a *strong generator* $\mathcal{L}$ (also named *infinitesimal* generator), which can be thought of as being the derivative of the semigroup at the initial time ($t = 0$). Let $\mathcal{D}(\mathcal{L}) \subseteq C_b(\mathcal{R})$ be the set of functions f, for which the following operation exists:

$$\lim_{t \downarrow 0} \frac{1}{t}(P_t f - f).$$

Denote the above limit as $\mathcal{L}f$, where the convergence is intended to be in the sup norm. Notice that such an operator is defined by specifying the above limit, as well as its domain $\mathcal{D}(\mathcal{L})$. The following holds [Davis, 1993, Prop. 14.13]:

Proposition 3. *For $f \in \mathcal{D}(\mathcal{L})$, define the real-valued process*

$$M_t = f(\mathbf{s}(t)) - f(s_0) - \int_0^t \mathcal{L}f(\mathbf{s}(u))du.$$

Then, for any $\mathbf{s}(0) = s_0 \in \mathcal{R}$, M_t is a martingale on $(\Omega, \mathscr{F}, (\mathscr{F}_t)_{t \in \mathbb{R}^+}, \mathcal{P})$. □

For details on the concept of martingale, see page 21 and refer to [Varaiya, 1975; Billingsley, 1995; Borkar, 1995; Durrett, 2004].

The knowledge of a process generator often allows one to characterize its associated stochastic process. As it shall be argued, it is also beneficial in deriving conclusions on convergence properties of the process. This should suggest that it is desirable to find an explicit form for the generator associated with a stochastic process which is a solution of a stochastic model. Most of the computations involving real-valued functions $f \in C_b(\mathcal{R})$ of the stochastic process under study involve a relation known as *Dynkin formula*. This formula says that, for any $f \in \mathcal{D}(\mathcal{L}), \forall t \geq 0$,

$$E_s f(\mathbf{s}(t)) = f(s) + E_s \int_0^t \mathcal{L}f(\mathbf{s}(u))du. \tag{1.5}$$

This formula holds a certain importance, because of the applicative value of computing the expectation of functions of the stochastic process. We shall come back to this tenet in sections 1.4.3 and 2.2.1. Notice also that the Dynkin formula connects between SDEs (and, more generally, the theory of stochastic processes) and PDEs.

For the sake of generality, we may be interested in relaxing the claim in Proposition 3. Accordingly, the following entity is introduced.

Definition 7 (Extended Generator)**.** *Consider the functions $f \in \mathcal{D}(\mathcal{L}) \subseteq C_b(\mathcal{R})$ that verify the following: there exists a Borel measurable function $g : \mathcal{R} \to \mathbb{R}, g \in B(\mathcal{R})$, such that, for any $t \geq 0, \mathbf{s}(0) = s_0 \in \mathcal{R}$ the quantity*

$$M_t = f(\mathbf{s}(t)) - f(s_0) - \int_0^t g(\mathbf{s}(u))du \tag{1.6}$$

is a local *martingale (see page 21). Let us then define the* extended generator *of the process $\mathbf{s}(t)$ to be the operator $\mathcal{L}^e$ such that that $g = \mathcal{L}^e f$, with domain $\mathcal{D}(\mathcal{L}^e)$ to be the set of functions $f \in C_b(\mathcal{R})$ such that there exists a $g \in B(\mathcal{R})$ that verifies the relation (1.6).* □

The above definition is justified in terms of Proposition 3, because the concept of local martingale is weaker and easier to characterize than that of martingale [Varaiya, 1975; Billingsley, 1995; Borkar, 1995; Durrett, 2004]. In the rest of the dissertation, let us say that a càdlàg Markov process $\mathbf{s}(t)$, with values on $\mathcal{S}$, is a *solution* of the [local] martingale problem $(\mathcal{L}, s_0)$ $[(g, s_0)]$, if for any $f \in \mathcal{D}(\mathcal{L})$, M_t in Proposition 3 [M_t in (1.6)] is a martingale [a local martingale].

Consider a real-valued function f of the hybrid state space $\mathcal{S}, f : \mathcal{S} \to \mathbb{R}$.

Assumption 2. *Assume f is a class $C_b^2(\mathcal{S})$ function, that is a real-valued, bounded, twice continuously differentiable function, defined on $\mathcal{S}$.* □

The extended generator for the process associated with the GSHS model in Definition 5 is derived along the ideas developed in the seminal work of [Davis, 1993], and the extension to the diffusion case presented in [Bujorianu and Lygeros, 2004b].

To the stochastic process associated with the GSHS $\mathscr{S}_g$, let us associate the extended generator $\mathcal{L}_g$ as follows.

Definition 8 (Extended Generator of $\mathscr{S}_g$)**.** *Assume $\mathscr{S}_g$ verifies Assumptions 1.6 and 1.9. The extended generator $\mathcal{L}_g : \mathcal{D}(\mathcal{L}_g) \to B_b(\mathcal{S})$, associated with the solution of $\mathscr{S}_g$, is an operator acting on real functions f, with domain $\mathcal{D}(\mathcal{L}_g)$ containing all the f*

verifying Assumption 2. For $s = (q, x) \in \mathcal{S}$ and $f \in \mathcal{D}(\mathcal{L}_g)$, $\mathcal{L}_g f$ is given by

$$\begin{aligned}
\mathcal{L}_g f(s) &= \mathcal{L}_g^d f(s) + \sum_{q' \neq q \in \mathcal{Q}} \lambda_{qq'}(s) \int_{\mathbb{R}^{n(q')}} (f(s' = (q', x')) - f(s)) R^{\Lambda}(ds' = (q', dx'), s), \\
\mathcal{L}_g^d f(s) &= \sum_{i \in \mathcal{Q}} \frac{\partial f(q,x)}{\partial x_i} a_i(q,x) + \frac{1}{2} \sum_{j,k \in \mathcal{Q}} \sum_{l \in \mathcal{Q}} b_{jl}(q,x) b_{kl}(q,x) G^f(q,x), \\
G^f(q,x) &= \left[g_{ij}^f(q,x) \right]_{i,j \in \mathcal{Q}}, \quad \text{where} \quad g_{ij}^f(q,x) = \frac{\partial^2 f(q,x)}{\partial x_j \partial x_k}, \quad s \in \mathcal{S} \backslash \Gamma. \\
f(s) &= \int_{\mathcal{S}} f(s') R^{\Gamma}(ds', s) = \sum_{q' \neq q \in \mathcal{Q}} \int_{\mathbb{R}^{n(q')}} f(q', x') R^{\Gamma}((q', dx'), s), \quad s \in \Gamma.
\end{aligned}$$

□

Remark 5. *The above definition should more properly be a theorem, whereby it is proved that an operator so defined, verifies the local martingale condition in Definition 7 on functions belonging to its domain. However, the proof of such a statement closely follows that in [Bujorianu and Lygeros, 2004b, Theorem 2], which in turn strictly adheres to [Davis, 1993, 26.14]. Notice that, unlike those sources, no bounded-variation of the expected value of the functions is assumed. This is thanks to the simplifying boundedness hypothesis for f, as in Assumption 2, and on Assumption 1.9 for the GSHS model, which excludes Zeno behaviors (that is, an infinite number of transitions in a finite time interval).*

In the formulas for $\mathcal{L}_g$, the quantity $\mathcal{L}_g^d$ includes the contribution of the continuous part of the HS as it encompasses operations on its drift and diffusion terms. The second summand in the definition of $\mathcal{L}_g$ describes the influence of the transition intensities, and their related reset kernels. These terms act on points of the hybrid state space that are away from the guard set. Finally, the boundary condition at the bottom line accounts for the resets due to the spatial constraints, and in fact acts on hybrid points belonging to the guard set Γ. Notice that this last condition effectively restricts the domain of the operator $\mathcal{L}_g$. □

Remark 6 (Special Cases)**.** *If the diffusion term B in $\mathscr{S}_g$ is neglected, the extended generator $\mathcal{L}_g$ is included in that of PDMPs, as formally derived in [Davis, 1993, 26.14]. More precisely, their structures coincide, but the PDMP's would have a larger domain of definition (that of functions of class $C_b^1(\mathcal{S})$).*

Assume that the GHSH $\mathcal{S}_g$ has no guards, and (deterministic) identity reset maps (in other words, $\Gamma = \varnothing$ and $\forall j \neq i, i \in \mathcal{Q}, R^{\Gamma}(j, \{x\}, (i, x)) = R^{\Lambda}(j, \{x\}, (i, x)) = 1$). We obtain the framework known in the literature as "switching diffusions." The extended generator from Definition 8 for this special case coincides with that derived in the literature, for instance in [Ghosh et al., 1992], modulo the neglect of the two reset conditions, as well as of the boundary restriction.

Furthermore, to make an intuitive parallel, in the purely deterministic (no diffusion terms, no transition intensities, nor discrete events and corresponding reset kernels) and dynamical (single-domain) case, notice that the extended generator coincides with the Lie derivative of the function f, taken along the corresponding vector field [Sastry, 1999]. □

1.3 Elimination of Guards

The objective of this section is to approximate the GSHS model described in section 1.2.1 with a new SHS, where proper transition intensities are replaced to the spatial guards. Intuitively, the idea is to "substitute" the events due to the intersection of the trajectory with the guard set with "random events," which are sampled from proper probability distributions. The new model will then be characterized exclusively by jumps that are elicited by transition intensities. This procedure is a generalization and a formal refinement of the technique first proposed in [Abate *et al.*, 2005], which was in turn inspired by a preliminary comment in [Hespanha, 2004].

1.3.1 Approximation Procedure

Consider the GSHS system $\mathscr{S}_g$ and its guard set $\Gamma = \{\gamma_{i,j}\} \subset D_i, i \in \mathcal{Q}$. Assume that the sets $\gamma_{i,j} \subset D_i$ can be expressed as zero sublevel sets of continuous functions $h_{i,j} : \mathbb{R}^{n(i)} \to \mathbb{R}$, properly defined on the continuous part of the domain D_i:

$$\gamma_{i,j} = \{x \in \mathbb{R}^{n(i)} : h_{i,j}(x) \leq 0\}.$$

Pick a number $\delta \geq 0$, and leveraging the continuity of the functions $h_{i,j}$, introduce the sets[8]

$$\gamma_{i,j}^{-\delta} = \{x \in D_i : h_{i,j}(x) \leq -\delta\} \subseteq \gamma_{i,j} \subseteq \gamma_{i,j}^{\delta} = \{x \in D_i : h_{i,j}(x) \leq \delta\}.$$

Given a point $z \in \mathcal{S}$ and a set $A \subseteq \mathcal{S}$, let $d(z, A) = \inf_{y \in A} ||z - y||$ denote the distance between z and A. Furthermore, for any $j \neq i, j \in \mathcal{Q}$, introduce the set of functions $\lambda_{i,j}^{\delta} : D_i \to \mathbb{R}^+$

$$\lambda_{i,j}^{\delta}(x) = \begin{cases} 0, & x \in D_i \setminus \gamma_{i,j}^{\delta} \\ \left(\frac{1}{d(x,\gamma_{i,j}^{-\delta})} - \frac{1}{\sup_{\{y:h_{i,j}(y)=\delta\}} d(y,\gamma_{i,j}^{-\delta})}\right) \wedge \left(\frac{1}{\sup_{\{y:h_{i,j}(y)=0\}} d(y,\gamma_{i,j}^{-\delta})}\right), & x \in \gamma_{i,j}^{\delta} \end{cases} \tag{1.7}$$

[8] Let us assume that there exists a function $h_{i,j}$ such that, for small enough $\delta > 0$, both $\gamma_{i,j}^{\delta}, \gamma_{i,j}^{-\delta} \neq \varnothing$. The first condition happens if the interior of a domain is non-trivial. The second if the guard is not strictly inside the corresponding domain and it does not have a trivial volume.

In the above expression, the operation $a \wedge b = \min\{a, b\}$ has been used. Associated with this set of intensity functions is a $\Lambda^\delta : \mathcal{S} \times \mathcal{Q} \to \mathbb{R}^+$. For any $0 < \delta < \infty$ the function Λ^δ is upper bounded. In the limit as $\delta \to 0$, the following set of intensity functions is obtained

$$\lambda^0_{i,j}(x) = \begin{cases} 0, & x \in D_i \backslash \gamma_{i,j}, \\ +\infty, & x \in \gamma_{i,j}. \end{cases}$$

With the system $\mathscr{S}_g$ is associated a new stochastic hybrid system $\mathscr{S}_\delta$, which is made up of the elements of $\mathscr{S}_g$, except for the following:

- The spatial guards set is empty, $\Gamma = \varnothing$;
- The new domains are simply $D_i = \mathbb{R}^{n(i)}, \forall i \in \mathcal{Q}$;
- The new set of transition intensities Λ^δ is defined over the whole $\mathcal{S}$;
- The old set of transition intensities Λ, whose original domain of definition was $\mathcal{S} \setminus \Gamma$, are extended to the whole domain by setting their value to be equal to zero for the points inside the old guard set;
- The reset kernel R^Γ (for Λ^δ), which used to be defined over the domain of the old guard set Γ of $\mathscr{S}_g$, is extended to the whole $\mathcal{S}$ by introducing deterministic identity resets on the points in $\mathcal{S} \backslash \Gamma$;
- The reset kernel R^Λ (for Λ), which used to be defined just on the interior $\mathcal{S} \backslash \Gamma$ of the domains of $\mathscr{S}_g$, is extended to the whole $\mathcal{S}$ also by introducing deterministic identity resets on the points in Γ.

Let us remark that the absence of spatial guards implies that the events associated with a hybrid execution are random events exclusively due to the presence of the transition intensities (both the old and the newly defined ones). It can be formally shown that the evolution of the discrete component can be described by a non-homogeneous continuous-time Markov chain: section 1.4.1 further elaborates this point and leverages this feature.

It is worthwhile to add that the definition of the transition intensities in (1.7) is not unique. Assuming some simplified form for the guard set, [Abate *et al.*, 2005]

has come up with an alternative and less general definition. Furthermore, for the simple instance of the hybrid model of the *bouncing ball*, [Hespanha, 2004] proposes yet another set of parameter-dependent intensities. Regardless of the actual shape of the functions, it is their limiting properties that will draw the attention in the following.

1.3.2 Convergence Properties

Let us derive the extended generator for the stochastic hybrid process, which is a solution of $\mathscr{S}_\delta$. The form of this generator is based on results in the literature. [Davis, 1993] derives the extended generator for a PDMP. In [Jacod and Shiryaev, 1987; Ghosh *et al.*, 1992], the generator of a switched diffusion is reported. [Ghosh and Bagchi, 2004] extend this form to a switched diffusion with deterministic resets. The generator here derived is then more general than those cited in that the SHS has random resets associated with the discrete jumps. Consider again real-valued functions f, defined on the hybrid state space $\mathcal{S}$ of $\mathscr{S}_\delta$, $f : \mathcal{S} \to \mathbb{R}$.

Definition 9 (Extended Generator of $\mathscr{S}_\delta$)**.** *Assume that $\mathscr{S}_\delta$ verifies Assumptions 1.6 and 1.9.*[9] *The extended generator $\mathcal{L}_\delta$, associated with the solution of $\mathscr{S}_\delta$, is an operator acting on real-valued functions f, with domain $\mathcal{D}(\mathcal{L}_\delta)$ containing all the f verifying Assumption 2. For $s = (q, x) \in \mathcal{S}$, and for $f \in \mathcal{D}(\mathcal{L}_\delta)$, $\mathcal{L}_\delta f$ is defined as*

$$\begin{aligned}
\mathcal{L}_\delta f(s) &= \mathcal{L}_\delta^d f(s) + \sum_{q' \neq q \in \mathcal{Q}} \lambda_{qq'}(x) \int_{\mathbb{R}^{n(q')}} (f(s' = (q', x')) - f(s)) R^\Lambda(ds' = (q', dx'), s) \\
&\quad + \sum_{q' \neq q \in \mathcal{Q}} \lambda^\delta_{qq'}(x) \int_{\mathbb{R}^{n(q')}} (f(s' = (q', x')) - f(s)) R^\Gamma(ds' = (q', dx'), s), \\
\mathcal{L}_\delta^d f(s) &= \sum_{i \in \mathcal{Q}} \frac{\partial f(q,x)}{\partial x_i} a_i(q, x) + \frac{1}{2} \sum_{j,k \in \mathcal{Q}} \sum_{l \in \mathcal{Q}} b_{jl}(q, x) b_{kl}(q, x) G^f(q, x), \\
G^f(q, x) &= \left[g^f_{ij}(q, x) \right]_{i,j \in \mathcal{Q}}, \quad \text{where} \quad g^f_{ij}(q, x) = \frac{\partial^2 f(q,x)}{\partial x_j \partial x_k}.
\end{aligned}$$

□

Remark 7. *Notice the absence of the guard conditions of Definition 8. It is substi-*

[9] As anticipated on page 19, in section 1.3.3 it will be shown that, for the SHS model $\mathscr{S}_\delta$, 1.7 is sufficient for 1.9 to hold true.

tuted by the third additive term in the definition of $\mathcal{L}_\delta$, where the "artificial" transition intensities Λ^δ, as defined in equation (1.7), have been introduced. In principle, the absence of the guard condition in Definition 9 makes $\mathcal{D}(\mathcal{L}_g) \subseteq \mathcal{D}(\mathcal{L}_\delta)$, for $\delta > 0$. □

For the ease of notations, let us introduce a change of variables: $n = 1/\delta$. Assuming without effective loss of generality that $n \in \mathbb{N}$, in the following let us refer to the SHS $\mathscr{S}_n$, in place of $\mathscr{S}_\delta$. Let us denote with $\mathbf{s}(t)$ and $\mathbf{s}_n(t), t \geq 0$, the stochastic processes solutions of, respectively, the GSHS $\mathscr{S}_g$ and $\mathscr{S}_n$. Recall the definition of their extended generators, $\mathcal{L}_g$ in Definition 8 and $\mathcal{L}_n = \mathcal{L}_{1/\delta}$, given just above.

Let us formally show that, as $n \to \infty$, the sequence of stochastic processes $\{\mathbf{s}_n(t)\}_{n\geq 1}$ converges, in some sense, to $\mathbf{s}(t)$, for any $t \geq 0$. In other words, let us show that, at the limit, the solutions of the GSHS $\mathscr{S}_g$ and of the SHSs $\mathscr{S}_n$ are, in some sense, the same. In the following let us pin down the notion of convergence. The forthcoming concepts can be found in [Ethier and Kurtz, 1986; Jacod and Shiryaev, 1987], where a number of results on weak convergence of stochastic processes that will be used in the following are summarized.

Consider $\mathcal{C}_b^0(\mathcal{S})$, the space of real-valued, bounded and continuous functions f defined on the hybrid state space $\mathcal{S}$ and endowed with the sup norm, $\|f\| = \sup_{s\in\mathcal{S}} |f(s)|$. Let $\mathscr{P}(\mathcal{S})$ be the space of probability distributions over $\mathcal{S}$.

Definition 10 (Weak Convergence)**.** *A sequence $\{\mu_n\}_{n\geq 1} \subset \mathscr{P}(\mathcal{S})$ of probability distributions is said to* converge weakly *to $\mu \in \mathscr{P}(\mathcal{S})$ if*

$$\lim_{n\to\infty} \int_{\mathcal{S}} f d\mu_n = \int_{\mathcal{S}} f d\mu, \quad \forall f \in C_b^0(\mathcal{S}).$$

□

Consider a random variable X with values in $\mathcal{S}$. Denote its probability distribution, an element in $\mathscr{P}(\mathcal{S})$, as $\mathtt{dist}_X : \mathcal{B}(\mathcal{S}) \to [0,1] : \int_{\mathcal{S}} \mathtt{dist}_X(x)dx = 1$. If $A \in \mathcal{B}(\mathcal{S}), \mathtt{dist}_X(A) = \mathcal{P}(X \in A)$.

Definition 11 (Convergence in Distribution)**.** *A sequence of $\mathcal{S}$-valued random variables $\{X_n\}_{n\geq 1}$ is said to* converge in distribution *to an $\mathcal{S}$-valued random variable X if $\{\mathtt{dist}_{X_n}\}_{n\geq 1}$ converges weakly to $\mathtt{dist}_X$, that is*

$$\lim_{n\to\infty} \int_{\mathcal{S}} f(z)\mathtt{dist}_{X_n}(dz) = \int_{\mathcal{S}} f(z)\mathtt{dist}_X(dz).$$

Equivalently,

$$\lim_{n\to\infty} E[f(X_n)] = E[f(X)], \quad \forall f \in C_b^0(\mathcal{S}). \qquad \square$$

With the understanding of the differences between the two notions of convergence, we shall use the notations $\mu_n \Rightarrow \mu$ and $X_n \Rightarrow X$ respectively.

The concept of extended generator (section 1.2.2) can be useful in showing that a sequence of Markov processes converges to a given Markov process. Informally [Ethier and Kurtz, 1986], given a sequence of $\mathcal{S}$-valued processes $\{\mathbf{X}_n\}_{n\geq 1}$ and a process $\mathbf{X}$, endowed with extended generators $(A_n, \mathcal{D}(A_n))$ and $(A, \mathcal{D}(A))$ respectively, to prove that $\mathbf{X}_n \Rightarrow \mathbf{X}$ it is sufficient to show that for all functions $f \in \mathcal{D}(\mathcal{A})$, there exist $f_n \in \mathcal{D}(\mathcal{A}_n)$, such that $f_n \to f$ and $A_n f_n \to Af$.

Let us introduce the following condition:

Definition 12 (Compact Containment Condition (CCC))**.** *A sequence of stochastic processes $\{\mathbf{X}_n\}_{n\geq 1}$ on $\mathcal{S}$ is said to satisfy the* Compact Containment Condition *if for any $\epsilon > 0, N > 0$, there exists a compact set $K_{\epsilon,N} \subset \mathcal{S}$ such that*

$$\liminf_{n\uparrow\infty} \mathcal{P}\left[\mathbf{X}_n(t) \in K_{\epsilon,N}, \forall\, 0 \leq t \leq N\right] \geq 1 - \epsilon. \qquad \square$$

The following is verified:

Theorem 1. *Consider the GSHS $\mathscr{S}_g$ and the SHS $\mathscr{S}_n, \forall n$. Let $\mathbf{s}(t)$ and $\mathbf{s}_n(t), \forall n$, be the unique global solutions of these models, $t \geq 0$. Then, the stochastic processes $\mathbf{s}(t)$ and $\mathbf{s}_n(t)$ verify the CCC in Definition 12.* $\square$

Proof. The existence and uniqueness of (global) solutions for the hybrid models requires the validity of a number of assumptions on them (see Proposition 1). Among these, Assumption 1.3 requires boundedness of both the drift and the diffusion terms. In addition, Assumption 1.8 implies that the possible resets have a finite "range."

This is enough to argue for the absence of a finite escape time for the stochastic processes that are solution of the SHS $\mathscr{S}_g$ and $\mathscr{S}_n$. In other words, the solution does not diverge in finite time: $\forall 0 \leq N < \infty, \sup_{u\in[0,N]} \|\mathbf{s}_\cdot(u) - \mathbf{s}_\cdot(0)\| < \infty$. Because of this feature, at any time $N \geq 0$, it is always possible to find a set, in particular a

compact one, that contains the trajectories with probability one for the whole time interval $[0, N]$ and for any approximation parameter $n = 1/\delta$. The last formula may suggest a lower bound for the diameter of such compact set. This argument concludes the proof. □

In the rest of the dissertation, given a sequence of entities $\{c_n\}_{n\geq 1}$ and a c, let us denote with $\lim_n^\star c_n = c$ the conditions $\lim_{n\to\infty} c_n = c$ and $(\vee_n \|c_n\|) \vee \|c\| < \infty$, where again $\|\cdot\|$ is the supremum norm.

The following statement, drawn from [Xia, 1994, Theorem 4.4] and proven in generality there, describes a condition for the weak convergence of Markov Processes by the use of their extended generators. The result is based on some theory developed and contained in [Ethier and Kurtz, 1986; Jacod and Shiryaev, 1987].

Proposition 4 (Conditions for Weak Convergence)**.** *Let $(A, \mathcal{D}(A))$ and $(A_n, \mathcal{D}(A_n))_{n\geq 1}$ be the extended generators of the $\mathcal{S}$-valued stochastic processes $\mathbf{X}(t)$ and $(\mathbf{X}_n(t))_{n\geq 1}, t \geq 0$. The initial conditions for these processes are sampled respectively from π and $(\pi_n)_{n\geq 1}$ in $\mathscr{P}(\mathcal{S})$, all probability distributions on $\mathcal{S}$. Assume that the processes $(\mathbf{X}_n)_{n\geq 1}$ are solutions of the local martingale problem (see page 25) for $(A_n, \pi_n)_{n\geq 1}$, and that $\mathbf{X}$ is a unique solution of the local martingale problem for (A, π). Suppose further that:*

- $A \subset C_b^0(\mathcal{S}) \times C^0(\mathcal{S})$;
- *For all $f \in \mathcal{D}(A), \exists f_n \in \mathcal{D}(A_n), n \geq 1$, such that $\lim_n^\star f_n = f, \lim_n A_n f_n = Af$;*
- $\pi_n \Rightarrow \pi$.

If $(\mathbf{X}_n)$ satisfies the CCC in Definition 12 and $\mathcal{D}(A)$ is dense in $C_b^0(\mathcal{S})$ with respect to $\lim^\star$, then $\mathbf{X}_n \Rightarrow \mathbf{X}$, as $n \to \infty$. □

Notice that the generators A and $(A_n)_{n\geq 1}$ are not required to have the same domain. Also, recall that a process $\mathbf{X}$ is said to be a solution of the local martingale problem for a linear operator (A, π) if $\mathcal{P} \circ \mathbf{X}(0)^{-1} = \pi$, and for each $f \in \mathcal{D}(A)$, $f(\mathbf{X}(t)) - f(\mathbf{X}(0)) - \int_0^t Af(\mathbf{X}(s))ds$ is a local martingale, $\forall t \geq 0$. Clearly this property holds in particular, by Definition 8, for the extended generator of the process $\mathbf{X}$.

Similar arguments also hold for the sequence of processes $\{X_n\}_{n\geq 1}$.
The upcoming hypothesis can be related to Assumption 1, as well as Assumption 5.

Assumption 3. *Given a GSHS, as in Definition 5, assume that the probabilistic reset kernels $R^{\Gamma}(j,\cdot,(q,x))$ are continuous in x.* □

In order to complete the proof of the following statement 2, it is necessary to verify the following claim, the proof of which is, as of yet, still elusive.

Conjecture 1. *The local martingale problem for $(\mathcal{L}_g, \mathcal{D}(\mathcal{L}_g))$ is well posed, that is, it admits a* unique *solution.* □

The result in Proposition 4 can be applied to the (G)SHS setup, thus yielding the following theorem:

Theorem 2 (Weak Convergence of $\mathscr{S}_\delta$ to $\mathscr{S}_g$)**.** *Consider the SHS model $\mathscr{S}_\delta$, the GSHS model $\mathscr{S}_g$ under Assumption 3, and their associated solution processes $\mathbf{s}_\delta(t)$ and $\mathbf{s}(t), t \geq 0$, where $\mathbf{s}_\delta(0) = \mathbf{s}(0) = (q_0, x_0) \in \mathcal{S}$. Consider further their extended generators $(\mathcal{L}_\delta, \mathcal{D}(\mathcal{L}_\delta)), (\mathcal{L}_g, \mathcal{D}(\mathcal{L}_g))$ and assume Conjecture 1 is valid. As the approximation step $\delta \downarrow 0$, the solution of the SHS $\mathscr{S}_\delta$ weakly converges to that of the GSHS $\mathscr{S}_g$: $\mathbf{s}_\delta(t) \Rightarrow \mathbf{s}(t), \forall t \geq 0$.* □

Proof. It will be ordinately shown that the conditions in Proposition 4 hold true for the theorem under study. To begin with, let us relate the SHS models $\mathscr{S}_\delta$ with the associated $\mathscr{S}_n$, its corresponding solution $\mathbf{s}_n(t)$ and the extended generator $(\mathcal{L}_n, \mathcal{D}(\mathcal{L}_n))$, where it is safely assumed that $n \in \mathbb{N}$. Notice that the initial probability distributions associated with the processes $\mathbf{s}$ and $\mathbf{s}_n$ coincide, and are in fact concentrated on the hybrid point $(q_0, x_0) \in \mathcal{S}$.

As previously discussed, because of the boundary condition raised in Definition 8, for any $n \geq 1$, the generators $\mathcal{L}_n$ and $\mathcal{L}_g$ have different domains. Furthermore, $\mathcal{L}_n$ has a shape that differs from $\mathcal{L}_g$, in that there is a contribution from the newly created intensity functions Λ^δ and their associated reset maps.

As discussed in Remark 5 both extended generators $(\mathcal{L}_n, \mathcal{D}(\mathcal{L}_n))$ and $(\mathcal{L}_g, \mathcal{D}(\mathcal{L}_g))$ verify the local martingale property on functions f of their solution processes, and which belong to their domain of definition. The stochastic processes $s_n(t)$ and

$\mathbf{s}(t)$ are solutions of the corresponding local martingale problems $(\mathcal{L}_n, (q_0, x_0))$ and $(\mathcal{L}_g, (q_0, x_0))$ [Davis, 1993]. However, we have to raise in Conjecture 1 the claim that $\mathbf{s}(t)$ is a unique solution.

Notice that the deterministic initial condition (q_0, x_0) makes the third condition in Proposition 4 unnecessary.

Given the domain of the extended generators, which includes $f \in C_b^2(\mathcal{S})$, as well as the form of the generators, the first of the three conditions in Proposition 4 certainly holds for points $s \in \mathcal{S} \setminus \Gamma$. Furthermore, the continuity assumption on the reset kernel R^Γ in Assumption 3 enforces this condition on functions $f \in C_b^2(\mathcal{S})$, which verify the third condition in Definition 8, on points $s \in \Gamma$.

For any $n \geq 1$, let us introduce the functions $f_n : \bigcup_{q\in\mathcal{Q}}\{q\}\times\mathbb{R}^{n(q)} \to \mathbb{R}$, to be identical to $f : f_n = f, \forall s \in \bigcup_{q\in\mathcal{Q}}\{q\}\times\mathbb{R}^{n(q)}$. Then, we do not need to assume a novel sequence of functions $f_n \in \mathcal{D}(\mathcal{L}_n)$ such that $\lim_n^\star f_n = f$, as in the second point of Proposition 4. Let us then limit ourselves to proving that $\lim_n \mathcal{L}_n f = \mathcal{L}_g f$. In this limit the boundedness of the involved quantities is not required. Recalling the shape of the generators $\mathcal{L}_n f$ and $\mathcal{L}_g f$ in Definitions 8 and 9, notice that the only difference lies on the presence of the third additive term, on one side, and the guard set condition, on the other. The reader should realize that the extension of the domains of $\Lambda, R^\Gamma, R^\Lambda$ does not introduce additional, spurious behaviors to the dynamics of $\mathscr{S}_\delta$.

Let us focus on the asymptotics of Λ^δ. First, on the points inside the domains, $s \in \mathcal{S} \setminus \Gamma = \bigcup_{q\in\mathcal{Q}}\{q\} \times D_q \setminus \Gamma, \lim_{n\to\infty} \Lambda^\delta(s) = 0$. Secondly, on the points of the guard set, $s \in \Gamma$, the value of the barrier functions diverges. Then, the following limit

$$\lim_{n\to\infty} \sum_{q'\neq q\in\mathcal{Q}} \lambda_{qq'}^n(q,x) \left\{ \int_{\mathbb{R}^{n(q')}} f(s' = (q', x')) R^\Gamma(ds' = (q', dx'), (q, x)) - f(q, x) \right\} = 0$$

will be verified only if the second multiplicative term is equal to zero, which coincides with the boundary condition for $\mathcal{L}_g f$ in Definition 8:

$$f(q,x) = \int_{\mathbb{R}^{n(q')}} f(s' = (q', x')) R^\Gamma(ds' = (q', dx'), (q, x)).$$

As it has been discussed in Theorem 1, the sequence of processes $(\mathbf{s}_n(t))_{n\geq 1}$ satisfies the CCC in Definition 12.

Let us now show that $\mathcal{D}(\mathcal{L}_g)$ is dense in $C_b^0(\mathcal{S})$. By the Stone-Weierstrass Theorem [Stone, 1948], given any compact set K in a Haussdorff space (in particular, the metric space $\mathcal{S}$), $C^2(K)$ functions are dense over the set of $C^0(K)$ functions. In fact, they form an algebra that contains constant functions and separates points. This holds also for the functions $C_b^2(K)$ over the set $C_b^0(K)$. The claim can be extended to the set of $C_b^2(\mathcal{S})$ functions, as needed in Proposition 4. First, notice that $\mathcal{S}$ is *locally* compact, as Euclidean spaces are. Then, refer to the extensions of the Stone-Weierstrass theorem to locally compact spaces, as discussed in [Nel, 1968; Arens, 1949], which can be leveraged in the framework of [Xia, 1994] that we are exploiting in Proposition 4.

It has been shown that all the conditions in Proposition 4 are verified. The claim of the theorem follows. □

One alternative way to go about proving such a convergence result is by leveraging the infinitesimal generator of the process. There is a body of literature on the subject, nicely summarized in [Ethier and Kurtz, 1986; Jacod and Shiryaev, 1987]. These results go through concepts, such that of a *core*[10] of an operator, which is fundamentally related to the "density" condition raised in Proposition 4, or that of *tightness*, which connects with the CCC. However, it appears difficult to derive a form for the infinitesimal generator of a GSHS ([Ghosh *et al.*, 1992] did this for the more restricted case of switched diffusions). On a side note, the above result can be extended to the case where the cardinality of $\mathcal{Q}$ is infinite. This instance may create problems for the validity of the CCC in Definition 12. A solution to the problem, and a consequent extension of the result in Theorem 2, is based on a martingale approach, as discussed in [Bhatt and Karandikar, 1993].

1.3.3 Convenience of the new form: A Claim

In Remark 3, it has been stressed how, in the deterministic case, "switched systems" are fundamentally simpler than "hybrid ones." The take-away point is that the

[10] Let A be a closed linear operator. A subspace D of $\mathcal{D}(A)$ is said to be a *core* for A if the closure of the restriction of A to D is equal to A.

presence of spatial guards increases the complexity of the model, possibly disrupts properties, and introduces pathological behaviors.

It is the author's belief that this statement should be extended to the stochastic realm. This has further motivated the introduction of the approximated models in this section.

More precisely, let us formalize the following claim, that *"stochastic hybrid systems which are endowed exclusively with random jumping events (which can be due to the presence of transition intensities or which, more generally, are due to an arrival process) are simpler than those that present spatial guards as a possible event mechanism."* In this claim, the attribute *simpler* has to be intended as a term of comparison between the original and the approximated models in terms of their structures, their behaviors, the properties that can be proven and the ease in simulating them.

A special instance of the above proposition is represented by the approximation of a deterministic HS, as in Definition 1, by a PDMP [Davis, 1993] endowed with transition intensities, but no guards.

This section is aimed at offering some supporting evidence in favor of the above claim. Often the argument will hinge on computational case studies which, for the sake of clarity, will be applied to the aforementioned special instance. A general statement for the above tenet is certainly not claimed here, but this work intends to reinforce the argument by offering the description of the many instances where this statement appears to be valid. The underlying message is that an approach such as that in section 1.3.1 may be exploited in the analysis and computation of models as comprehensive as the GHSH.

Absence of Zeno. Recall the concept of Zeno behavior, given in Definition 4. The following condition recaptures that in Assumption 1.9 on page 19, which was then used for proving the existence, uniqueness, and the Markov property of the solution process for the GSHS $\mathscr{S}_g$, and which will also play a role in the ensuing section.

Assumption 4. *Let us assume the transition intensities are bounded on their domain of definition, that is*

$$\exists 0 \leq \mathfrak{M} < \infty : \forall q' \in \mathcal{Q}, \forall s = (q, x) \in \mathcal{S}, \lambda_{qq'}(x) < \mathfrak{M}. \qquad \square$$

Theorem 3 (Absence of Zeno)**.** *Given a SHS with no spatial guards ($\Gamma = \varnothing$) (for instance $\mathscr{S}_n$ obtained via the approximation procedure in section 1.3.1), if Assumption 4 holds, then Zeno behaviors are ruled out, with probability one.* □

Proof. The survivor function in equation 1.4 on page 21 can be lower bounded by $F(t, s_0) \geq e^{-\mathfrak{M}t}$. The term on the right hand-side describes the probability that there is no jump in the interval $[0, t]$, according to an exponential distribution that may be associated with a Poisson arrival process with rate $\mathfrak{M}$. Then, the expected number of arrivals on $[0, t]$ is simply $\mathfrak{M}t$, hence it is finite for any finite time interval. It follows that the probability of an infinite number of arrival in a finite time interval is almost surely equal to zero. □

Remark 8. *[Abate et al., 2005, Theorem 1] proved a similar result for the special case mentioned above: the approximation of the dynamics of a deterministic HS with a PDMP [Davis, 1993], where the guard set is properly substituted by newly introduced transition intensities.*
As we discussed in a footnote of Assumption 1, the work in [Davis, 1993, Assumption 24.4, Proposition 24.6] derived explicit sufficient conditions to almost surely exclude Zeno behaviors, or the presence of finite "escape time." A comparison of those conditions with that in Assumption 4 will suggest the simplicity of the latter.
More generally, given any HS (or a GHSH) with guards, the application of an approximating procedure such as that in 1.3.1 yields a model which, upon verifying the condition in Assumption 4, is not Zeno. Let us stress that this assumption is practically *verified for simulation purposes.* □

Example 3 (Two Water Tanks Hybrid System, from [Abate *et al.*, 2005])**.** *The two water tanks hybrid system [Zhang et al., 2001; Abate et al., 2005; Ames et al., 2005] is a simple and well studied hybrid model of a physical system, which can present Zeno behavior among the allowed executions, depending on the value of some model parameters. Regardless of how one intends to regard and handle the concept of Zeno, it is of interest to be able to simulate such trajectories, without any possible disruptive outcome.*

The system represents the dynamics of the level of a liquid (say, water) in two tanks, set up in a parallel configuration. Both tanks, which allow for some liquid

outflow, may accept more water when filled up by a single hose. This refill mechanism is forced to switch between any of the two tanks whenever the corresponding level of fluid goes below a certain constant level (threshold). The device can be modeled as a deterministic HS. More precisely, let us introduce a set of two modes, $Q = \{1, 2\}$*, each of which is associated with a domain, the real line,* $D_1 = D_2 = \mathbb{R}^+$*. The domains are related via two symmetric edges,* $E = \{e_1 = (1, 2), e_2 = (2, 1)\}$*. A mode describes the active status of the corresponding tank, which is baing filled up. Accordingly, the other tank will instead be releasing liquid. The dynamics are described by the following two vector fields, acting on the corresponding domain and with values in* $\mathbb{R}^+$*:*

$$f_1(x) = \begin{pmatrix} -v_2 \\ w - v_1 \end{pmatrix}; \qquad f_2(x) = \begin{pmatrix} w - v_2 \\ -v_1 \end{pmatrix}.$$

Here $x = (x_1, x_2)^T$ *is the level of the water in the two tanks,* w *is the inflow,* v_1, v_2 *are the two outflows. The guards are two threshold levels (say* $\bar{x}_1, \bar{x}_2 \geq 0$*), and both reset maps are the identity (i.e., it is assumed that the switch of the external hose happens instantaneously).*

As said, the classical deterministic mathematical model of this system does not account for delays in the physical switch of the hose from one tank to the other one. It can be formally shown that this implies, for certain configurations of the parameters of the system, the existence of Zeno behavior: in particular, for the following case

$$\max\{v_1, v_2\} < w < v_1 + v_2.$$

In words, this happens in the case when the inflow is less than or equal to the sum of the two outflows. In this case, it is possible to calculate explicitly the exact value of the (finite) time when Zeno occurs (Zeno time t_∞*). In Fig. 1.3, we plotted first a 3D simulation of this system—under conditions that configure the system to be Zeno—with the classic deterministic event detection [Shampine and Thompson, 2000]. As expected, the number of switches of a trajectory diverges while approaching Zeno time.*

A proposed solution involves the substitution of the aforementioned constant thresholds $(\bar{x}_1, \bar{x}_2)$ *with barrier functions, whose values rapidly diverges in a neighborhood of the thresholds, while being close to zero inside the domains. The particular shape*

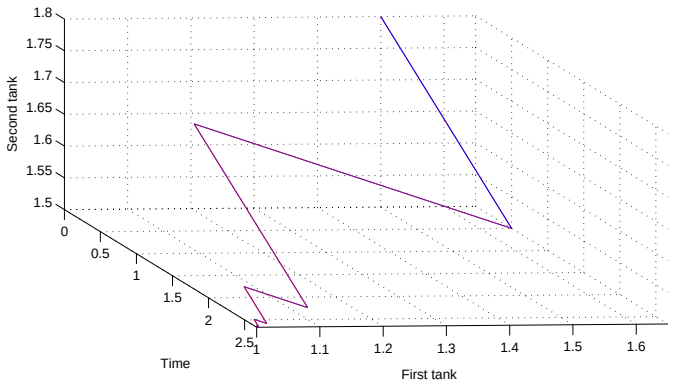

Figure 1.3: Three dimensional view of the two water tanks hybrid system.

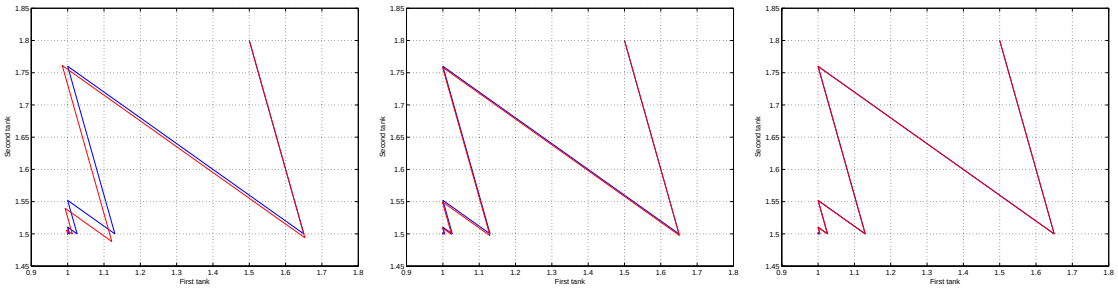

Figure 1.4: Top view of the two water tanks hybrid system. The model has been simulated with the classical deterministic event detection software available in MATLAB. After approximating the threshold via properly defined stochastic barrier functions, the simulations have been run with increasing precision by tuning a parameter for the intensity barriers.

of these transition intensities has been inspired by [Hespanha, 2004]. Their steepness depends on controllable parameters.

In the series of plots in Fig. 1.4, it is shown how the stochastic approximation is able to detect the discrete event with qualitatively as much accuracy as the deterministic counterpart. However, unlike this case, the trajectories are defined for any positive time interval. In other words, there is no finite Zeno time, and–from a computational perspective—the simulation does not stop. □

Deadlock Behaviors. Recently, the contribution in [Abate *et al.*, 2006b] proposes a formal definition of the concept of *deadlock* for a general class of (controlled) deterministic HS. It is argued that the concept arises because of 1. the composition of

different hybrid systems; and 2. the composition of *specifications* defined on trajectories belonging to their corresponding domains. In the categorization of the possible deadlock behaviors, it is clear that the vast majority of them is due to the interconnection of the guard sets pertaining to different HS, which yields Zeno or blocking behaviors. The work suggests that an approximation—along the lines of that of section 1.3.1—of the original deterministic HS with a probabilistic one may solve this issue. The reader is invited to refer to [Abate *et al.*, 2006b] for more details on the topic.

Hybrid System Composition. The issue of hybrid systems composition has recently drawn some attention [Tabuada *et al.*, 2004; Abate *et al.*, 2006b], yet it has not been completely sorted out and understood. However, it is clear that one of the big challenges in composing HS is to understand, in general, how the dynamics are affected, or preserved. This is in contrast with the ease in composing purely dynamical systems: the composition is in this case simply intended as a sharing of signals, and the system dynamics as a black-box that is unnecessary to fully characterize. This input-output approach [Lynch *et al.*, 2001; Lynch *et al.*, 2003] has in part shed some light on how to handle the problem. Still, a general solution does not appear to be within reach. Especially, the presence of guards seems to allow for a wealth of possible pathological behaviors (deadlock is only one of them).

On the contrary, the composition of HS with no guard set (in particular, SHS obtained according to the procedure in section 1.3.1, or—in the deterministic case—switched systems) is semantically definable in a straightforward way and does not involve the possible introduction of pathological or unexpected behaviors. The number of intensity functions can be easily manipulated, provided the assumptions on their continuity hold valid. The author is interested in further pursuing this avenue in his future work.

Simulations and Event Detection. As discussed in [Abate *et al.*, 2005] for the deterministic HS case, the issue of event detection can be problematic. This is true for particular interconnections of the guards of the system, which may yield pathological behaviors with respect to a subset of the allowed initial conditions. This issue is well

discussed in the work [Esposito *et al.*, 2001], which is based on seminal work on the topic by [Shampine and Thompson, 2000]. The use of a stochastic approximation may allow one to approach the problem in a much more "controlled" way: by tuning the steepness parameter of the transition intensity, it is possible to tune the event detection precision, regardless of possible pathological behaviors of the dynamics.

This is even truer for GHSH: while the event detection problem for diffusion dynamics with respect to a spatial guard is a rather hard problem, the corresponding issue of arrival generation for intensity-based switchings offers easier conditions (sampling from a probability distribution) for the simulation of the probabilistic dynamics. Refer also to section 1.4.1 and the "thinning procedure" there introduced.

Topological Properties. Recent research on deterministic HS [Simic *et al.*, 2005; Ames and Sastry, 2005a] has tried to investigate and understand their topological properties. The use of refined approaches borrowed from differential geometry, homology theory and category theory has shed some light on the issue, albeit at the expense of the complicated mathematics involved in the process. The SHS framework defined in this chapter, on the contrary, offers a rather simple and direct understanding of the above topic. In particular, notions of metric, distance and topology are straightforwardly introduced in section 1.2.1. It is the author's belief that the reason stems from the absence of spatial guards and their substitution with easy-to-handle transition intensities.

Explicit Dynamics. A further convenience coming out of models with no guards is the possibility of expressing their dynamics explicitly via a set of random measures. This is also connected to the applicability of time sampling techniques to the processes under study. These concepts will be further elaborated in section 1.4.

Use of the Process Generator

The previous section has argued for the structural and analytical simplicity of a certain class of stochastic hybrid systems. The main argument has been that the elimination—or proper substitution with approximated entities—of the spatial guards is associated with analytical and computational benefits. One instance where this

structural simplicity is evident is in the shape of the extended generator associated with a solution of the model. We have derived it in Definition 9 and compared it to the one for the GSHS in Definition 8. We have already stressed the importance of this entity, both theoretically (characterization of the process and of its convergence properties) and computationally (Dynkin formula, in equation (1.5)).

In the following let us suggest that the use of the (extended) generator may be useful in a number of other instances, which would be otherwise harder to analyze directly via the solution process, or the semigroup (when possible to express it explicitly).

Delay Systems. The Systems Theory literature has broadly studied dynamical models with delays [Gyory and Ladas, 1991]. These models naturally describe systems in which at least part of the signals experience a time lag, often due to a non-negligible transmission or processing slowdowns.
In particular, the use of such models in population dynamics, or more generally in the biological arena, is known [Gopalsamy, 1992] and helps in describing heterogeneous systems with different time scales. The recent contribution in [Lygeros *et al.*, 2006] extends this modeling effort to the stochastic hybrid framework, by setting up a model that describes a delayed, noisy and multidimensional version of the Lotka-Volterra population dynamics.

The continuous dynamics describe an n-dimensional stochastic delayed differential equation for the population level of each of the the n species. The parameters of the system depend on the state of a finite discrete-valued process, which changes according to the discrete dynamics. There is no reset upon transitioning between the discrete modes.

With the possible exclusion of the delayed term, this class of systems belongs to that of switching diffusions, and hence is a subset of the SHS given in section 1.3.1. Denote this new class of SHS as $\mathscr{S}_\tau$. This class of systems appears to be difficult to analyze with the classical approach on delayed differential equations, that is by introducing a set of generator functions, defined on the interval $[-\tau, 0)$, and looking at the associated infinite-dimensional system [Gyory and Ladas, 1991]. Instead, this model turns out to be quite easily characterizable via its extended generator.

With the exception of a newly-introduced spurious variable, which accounts for the delayed signal, the extended operator has an almost identical form as the one in Definition 9, and works on the same domain of definition. Refer to [Lygeros *et al.*, 2006] for the form of the dynamics and of the generator, as well as for an analysis of the model.

Polynomial Systems. The work in [Hespanha, 2005] introduces a SHS model which has polynomial components in the continuous dynamics, (deterministic) reset maps and transition intensities. It is there argued that the expression of the extended generator as a polynomial allows one to turn polynomial functions of the process solution of the model into other polynomials, thus inducing an algebraic structure. Known "moment closure" methods leverage this feature of the extended generator and allow one to approximate higher-order moments of the process with controllable accuracy. This has clear repercussions on computations that may be performed on the model.

1.4 Time Discretization

In this section, with reference to the approximated models obtained by the procedure explained in section 1.3.1, let us introduce an alternative expression of its dynamics in section 1.4.1. Leveraging this representation, let us then introduce a first-order time discretization scheme in section 1.4.2. By raising some continuity hypotheses on the elements of the SHS, the weak convergence property of the discrete-time process to the original continuous-time one is outlined in section 1.4.3.

This procedure is interesting in light of the objectives of this work. It is upon the obtained discrete-time process (properly enhanced with controls) that the concepts in chapter 2 shall be developed.

1.4.1 Hybrid Dynamics through Random Measures

Let us consider a SHS, as given in Definition 5 but with no guard set, and let this model be autonomous (with no controls). The event times are exclusively due to a random arrival process, their distribution having the shape in (1.4), with exclusive dependence on the set of state-dependent transition intensities. To reflect this feature, denote this SHS with $\mathscr{S}_\lambda$. The structural and analytical conveniences of this hybrid model has been previously discussed (in section 1.3.3). Among these qualities, we hinted at the fact that its dynamics (and consequently the semantics of the model) can be synthetically expressed via a set of (dynamical) relations, rather than by a scheme, as in Algorithm 6.

Consider the continuous-time Markov process $\mathbf{s}(t) = (\mathbf{q}(t), \mathbf{x}(t)), t \geq 0$, unique solution of $\mathscr{S}_\lambda$. For the sake of simplicity, let us assume that each domain has the same dimension n, that is $\forall q \in \mathcal{Q}, n(q) = n$. Also, let us consider a finite number of modes and denote them as $\mathcal{Q} = \{1, 2, \ldots, M\}, M = \texttt{card}(\mathcal{Q}) < \infty$. In other words, the process $\mathbf{s}(t)$ takes values in $\mathcal{S} = \mathcal{Q} \times (\mathbb{R}^n)$, the hybrid state space.

In the remainder of the section, let us show that the process $\mathbf{s}(t)$ can be regarded as a (marked-)point process, and that we can express its dynamics, and by large that of the stochastic hybrid system $\mathscr{S}_\lambda$, by the use of a *random measure* $\mu(dz, dt)$, to be formally introduced briefly. This measure characterizes the arrivals in time that depend on the transition intensities, and has marks on the hybrid state space $\mathcal{S}$

according to the reset kernels in $\mathscr{S}_\lambda$.

For the sake of reference, the results are based on related work in [Ghosh *et al.*, 1992; Glasserman and Merener, 2003; Glasserman and Merener, 2004; Krystul and Bagchi, 2004; Prandini and Hu, 2006a]. The expression that shall be obtained is to be in fact considered more general than those in the referenced articles. In particular, [Ghosh *et al.*, 1992] and [Prandini and Hu, 2006a] consider only switching diffusions (that is, with no reset of the continuous component upon a jumping event); [Glasserman and Merener, 2003; Glasserman and Merener, 2004] focus on jumping diffusions, that is processes described by a marked point process that is made up by a single-mode continuous diffusion component and a random measure with state-dependent transition intensities which probabilistically reset the trajectory. Finally, [Ghosh and Bagchi, 2004; Krystul and Bagchi, 2004] consider SHS that are similar to $\mathscr{S}_\lambda$, but with deterministic reset functions. Seminal work on the connections between SDEs and point processes is to be found in [Marcus, 1978], while [Jacod and Shiryaev, 1987, chapter II] contains a detailed presentation of the theory of random measures within the general framework of semimartingales.

In the ensuing work, let us first show that it is possible to symbolically describe the dynamics of $\mathscr{S}_\lambda$ by the following pair of equations, $\forall t \geq 0$:

$$d\mathbf{x}(t) = a(\mathbf{q}(t^-), \mathbf{x}(t^-))dt + b(\mathbf{q}(t^-), \mathbf{x}(t^-))d\mathbf{w}(t) + \int_{\mathcal{S}} \mu(dz, dt); \tag{1.8}$$

$$\mathcal{P}(\mathbf{q}(t+dt) = j | \mathbf{q}(t) = i, \mathbf{x}(s), \mathbf{q}(s), s \leq t) = \lambda_{ij}(\mathbf{x}(t))dt + o(dt), \tag{1.9}$$

$\forall j \in \mathcal{Q}, j \neq i$, where the first equation describes the continuous component, and the second the discrete one. Let us recall that in the second equation (1.9) the intensities are introduced, at time $t \geq 0$ and for $\mathbf{q}(t) = i$, as $\lambda_{ij}(\mathbf{x}(t)) = \lambda((i, \mathbf{x}(t)), j)$.

We already discussed (see Proposition 2 in section 1.2.1) the Markov property of the solution of the SHS $\mathscr{S}_\lambda$. This allows one to simplify the left-hand side of (1.9) as follows

$$\mathcal{P}(\mathbf{q}(t+dt) = j | \mathbf{q}(t) = i, \mathbf{x}(s), \mathbf{q}(s), s \leq t) = \mathcal{P}(\mathbf{q}(t+dt) = j | \mathbf{q}(t) = i, \mathbf{x}(t)).$$

The first equation in (1.8) describes the continuous dynamics (drift and diffusion

parts), as well as the probabilistic jumps and resets (contribution of the random measure). The quantity $\mu(dz, dt)$ is a random measure which accounts for the resets in the continuous component of the trajectory, and contains information on the occurrence of the jumps. This term is supposed to be "active" only when a transition happens. However, the reset probability in our model $\mathscr{S}_\lambda$ depends on the second relation in (1.9), through the state-dependent arrival rates that characterize the jumps of the discrete component of the hybrid trajectory. Intuitively, because of the definition of the reset probability distributions, the mark space of this random measure is the hybrid state space $\mathcal{S}$.

The *compensator* ν associated with the random measure μ, also known as its *intensity*, is a predictable measure which, for any measurable $f : \mathbb{R}^+ \times \mathcal{S}$, makes the quantity

$$\int_0^t \int_{\mathcal{S}} f(s, z)\mu(ds, dz) - \int_0^t \int_{\mathcal{S}} f(s, z)\nu(ds, dz)$$

a martingale, for any $t \geq 0$ [Jacod and Shiryaev, 1987; Cont and Tankov, 2004]. For the SHS $\mathscr{S}_\lambda$, it is possible to express ν in a closed form: considering a node $q' \in \mathcal{Q}$ so that at time $t \geq 0$, $\mathbf{q}(t) = q'$, while at time $t^- < t, \mathbf{s}(t^-) = (\mathbf{q}(t^-), \mathbf{x}(t^-))$, the intensity of the measure $\mu(dz, dt)$ is

$$\nu(dz, \mathbf{q}(t), (\mathbf{q}(t^-), \mathbf{x}(t^-))) = \lambda((\mathbf{q}(t^-), \mathbf{x}(t^-)), q')R^\Lambda(q', z, (\mathbf{q}(t^-), \mathbf{x}(t^-)))dzdt. \quad (1.10)$$

Notice that, at time $t \geq 0$, the intensity is characterized by the product of two terms from the definition of $\mathscr{S}_\lambda$, an intensity λ and a kernel R^Λ (function of the variable z). It is a random quantity because it depends on a random signal (the hybrid solution of the SHS $\mathscr{S}_\lambda$). The first contribution in (1.10) accounts for the random event times (which hinge on the state-dependent jumping intensities), while the second term accounts for the probabilistic resets. Notice that these two terms are taken from the elements in the Definition 5 of GSHS. Let us further assume that ν is an absolutely continuous function of its variables z, t, which allows one not to worry about imposing conditions for the validity of the equality $\nu(dz, dt) = \nu(z, t)\ dz\ dt$.

The following hypothesis can be related to Assumption 1, and in part coincides with Assumption 4.

Assumption 5 (Continuity and Boundedness Properties)**.** *The elements of* Λ *and of* R^{Λ} *of the SHS* $\mathscr{S}_{\lambda}$ *are Lipschitz continuous functions of their continuous component* x*, that is of the continuous part of the hybrid state space.*
Let Assumption 4, on the boundedness of the transition intensities in Λ*, hold true.* □

Before further working on the dynamical relation in (1.8), where the random measure (its reset part) depends the specific mode that the hybrid process jumps to, let us first focus the attention on the dynamics in (1.9), which can instead be manipulated independently. With regards to the discrete component $\mathbf{q}(t)$, let us introduce a construction which is much like the one in [Ghosh *et al.*, 1992; Prandini and Hu, 2006a]. The main idea is that it is possible to express the dynamics of an (non-homogeneous) continuous-time Markov chain by a random measure. Intuitively, it is possible to associate a continuous-time Markov chain to the discrete dynamics of the component $\mathbf{q}(t)$.

Let us introduce a set of intervals of the real line $\Delta_{ij}(x) \in \mathbb{R}, x \in D_i = \mathbb{R}^n, \forall i, j \in \mathcal{Q}, i \neq j$, as follows:

$$
\begin{aligned}
\Delta_{12}(x) &= [0, \lambda_{12}(x)), \\
\Delta_{13}(x) &= [\lambda_{12}(x), \lambda_{12}(x) + \lambda_{13}(x)), \\
&\vdots \\
\Delta_{21}(x) &= \left[\sum_{n=2}^{M} \lambda_{1n}(x), \sum_{n=2}^{M} \lambda_{1n}(x) + \lambda_{21}(x) \right), \\
&\vdots \\
\Delta_{ij}(x) &= \left[\sum_{n=1}^{i-1} \sum_{m=1, m\neq n}^{M} \lambda_{nm}(x) + \sum_{m=1, m\neq i}^{j-1} \lambda_{im}(x), \sum_{n=1}^{i-1} \sum_{m=1, m\neq n}^{M} \lambda_{nm}(x) + \sum_{m=1, m\neq i}^{j} \lambda_{im}(x) \right).
\end{aligned}
$$

In other words, the general interval $\Delta_{ij}(x), i \neq j \in \mathcal{Q}$, has length $\lambda_{ij}(x), x \in D_i$. Let us recall Assumption 5 on $\lambda_{ij}(x) < \infty, \forall i, j \in \mathcal{Q}, x \in D_i$, which ensures that the above intervals are finite, given the finite cardinality of $\mathcal{Q}$. Let us introduce the "global" interval $\Delta(x) = \bigcup_{i\neq j, i,j\in\mathcal{Q}} \Delta_{ij}(x)$, which has length $\lambda_{\Delta}(x) = \sum_{i\neq j=1}^{M} \lambda_{ij}(x)$. Compute the (possibly non-unique) point $\widehat{x} = \arg\sup_{x\in\mathbb{R}^n} \lambda_{\Delta}(x)$, and define the value $\widehat{\lambda} = \lambda_{\Delta}(\widehat{x}) = \sup_{x\in R^n} \lambda_{\Delta}(x)$ and the interval $\widehat{\Delta} = \Delta(\widehat{x}) = [0, \widehat{\lambda})$. Notice that the

quantity $\widehat{\lambda}$ is related to an analogous one defined on page 21.

At this point, let us introduce two bounded and piecewise constant functions

$$\begin{aligned} l &: \mathcal{S} \times \mathbb{R}^+ \to \{0, \pm 1, \pm 2, \ldots, \pm(M-1)\}, \\ m &: \mathcal{S} \times \mathbb{R}^+ \to \{0, 1\} \end{aligned}$$

as follows:

$$\begin{aligned} l(i,x,z) &= \begin{cases} j-i, & \text{if } z \in \Delta_{ij}(x), \\ 0, & \text{else;} \end{cases} \\ m(i,x,z) &= 1 - \delta(l(i,x,z)) = \begin{cases} 1, & \text{if } \Delta_{(i-1)M}(x) \leq z < \Delta_{iM}(x), i \neq M \\ 1, & \text{if } \Delta_{(M-1)M}(x) \leq z < \Delta_{M(M-1)}(x), i = M \\ 0, & \text{else;} \end{cases} \end{aligned}$$

where $\delta(n), n \in \mathbb{Z}$, is the Kronecker delta function.
The dynamics of the discrete component $\mathbf{q}(t)$ of the hybrid trajectory can be expressed, at any time $t > 0$, by the use of a random measure as:

$$d\mathbf{q}(t) = \int_{\widehat{\Delta}} l(\mathbf{q}(t^-), \mathbf{x}(t^-), z) p(dz, dt), \tag{1.11}$$

where $p(dz, dt)$ is a Poisson random measure with intensity $\mathscr{U}(dz) \times \widehat{\lambda} dt$, where $\mathscr{U}$ is a uniform probability distribution over the interval $\widehat{\Delta}$. In other words, the random measure $p(dz, dt)$ characterizes random arrival times according to an homogeneous exponential distribution of parameter $\widehat{\lambda}$, and has marks that are sampled uniformly in the interval $\widehat{\Delta}$.

On the side, notice that the dynamics of $\mathbf{q}(t)$ could also have been expressed as:

$$d\mathbf{q}(t) = \int_{\mathbb{R}} l(\mathbf{q}(t^-), \mathbf{x}(t^-), z) \tilde{p}(dz, dt),$$

with $\tilde{p}(dz, dt)$ is a Poisson random measure with intensity $\texttt{Leb}(dz) \times dt$, where $\texttt{Leb}$ is the Lebesgue measure over the real line [Ghosh *et al.*, 1992]. For convenience reasons that will become clearer in the ensuing discussion, in the rest of the dissertation let us however stick with the expression in (1.11).

Based on the expression in (1.11) for the dynamics in (1.9), it is possible to move on to further elaborate the dynamics in (1.8). Introduce the following:

Assumption 6. *Assume that the following holds on the reset kernels R^Λ of $\mathscr{S}_\lambda$:*

$$\forall s = (q,x) \in \mathcal{S}, q' \neq q \in \mathcal{Q},\ \exists 0 < \lambda_0 < \infty, g \in \mathscr{P}(\mathcal{Q}\times\mathbb{R}^n), g : \mathcal{Q}\times\mathcal{B}(\mathbb{R}^n) \to [0,1],$$
$$s.t.\ R^\Lambda(q',z,s) \leq \lambda_0\ g(q',z),\ \textit{for any } z \in \mathbb{R}^n. \qquad \square$$

This assumption is not too stringent. It is possible to introduce, for instance, a real-valued function $\tilde{f} : \mathcal{Q}\times\mathbb{R}^n \to \mathbb{R}^+$ as

$$\tilde{f}(q',z) = \sup_{x\in D_q} \max_{q\neq q', q\in\mathcal{Q}} R^\Lambda(q',z,(q,x));$$

and define a function $\tilde{\xi} : \mathcal{Q} \to [1,M]$ such that, for any $q' \in \mathcal{Q}$, $\tilde{\xi}(q') = \int_{\mathbb{R}^n} \tilde{f}(q',dz)$. Then, define $g(q',z) = \frac{\tilde{f}(q',z)}{\tilde{\xi}(q')}$.
Furthermore, introduce the function $\bar{f} : \mathbb{R}^n \to \mathbb{R}^+$, as $\bar{f}(z) = \max_{q'\in\mathcal{Q}} \tilde{f}(q',z)$ and the value $\bar{\xi} = \max_{q'\in\mathcal{Q}} \tilde{\xi}(q')$. Finally, define $f : \mathbb{R}^n \to \mathbb{R}^+$ as $f(z) = \frac{\bar{f}(z)}{\bar{\xi}}$.

Following the approach in [Glasserman and Merener, 2003], define now a function $\theta : \mathbb{R}^n \times \mathbb{R}^+ \times \mathcal{Q} \times \mathcal{S} \to \{0,1\}$ as:

$$\theta(z,u,q',(q,x)) = \begin{cases} 1, & \text{if } u < \frac{R^\Lambda(q',z,(q,x))}{\bar{\xi}f(z)}, \\ 0, & \text{else.} \end{cases} \tag{1.12}$$

Consider a Poisson random measure $p(dz,du,dt)$ with mark space $\mathbb{R}^n \times [0,1]$, and with intensity $\nu_\zeta(dz,du,dt) = f(z)dz \times \mathscr{U}(du) \times \bar{\xi}dt$, where $\mathscr{U}$ is uniformly distributed in the interval $[0,1]$. This random measure has marks $z \in \mathbb{R}^n$, distributed as $f(z)$, with $u \in [0,1]$ according to $u \sim \mathscr{U}([0,1])$, and random arrivals with (homogeneous) rate $\bar{\xi}$. The uniform random variable u implements an acceptance/rejection mechanism, according to the value of the function θ. With each jump time of the Poisson random measure is associated a mark (z,u).
The likelihood of θ to be non zero is, at time $t \geq 0$, conditional on z, the mode $q' \in \mathcal{Q}$ and the value of the trajectory $(\mathbf{q}(t^-),\mathbf{x}(t^-))$, equal to $\frac{R^\Lambda(q',z,(\mathbf{q}(t^-),\mathbf{x}(t^-))}{\bar{\xi}f(z)}$. Thus, a

point process associated with a new random measure

$$\zeta(dz, dt) = \int_0^1 \theta(z, u, q', (\mathbf{q}(t^-), \mathbf{x}(t^-)))p(dz, du, dt) \tag{1.13}$$

for any $q' \neq \mathbf{q}(t^-) \in \mathcal{Q}$, has a point in $[t, t+dt)$ with mark $z \in \mathbb{R}^n$ with a probability given by $R^{\Lambda}(q', z, (\mathbf{q}(t^-), \mathbf{x}(t^-)))dt + o(dt)$, conditioned on $(\mathbf{q}(t^-), \mathbf{x}(t^-)) \in \mathcal{S}$.

Then, let us exploit the relation in (1.13) by introducing an additional Poisson random measure $\mu(dz, du, d\gamma, dt)$ with mark space $\mathbb{R}^n \times [0,1] \times \Gamma$, where $\Gamma = [0, \bar{\xi}\widehat{\lambda})$, and intensity $f(z) \times \mathscr{U}_1(du) \times \mathscr{U}_2(d\gamma) \times \bar{\xi}\widehat{\lambda}dt$. The first probability distribution $\mathscr{U}_1$ is uniformly sampled over the interval $[0,1]$, while the second $\mathscr{U}_2$ over the interval Γ. By exploiting this newly introduced Poisson random measure μ, as well as the functions θ and m, it is possible to express the dynamics of the continuous component of the hybrid trajectory at time $t > 0$ as:

$$\begin{aligned} d\mathbf{x}(t) &= a(\mathbf{q}(t^-), \mathbf{x}(t^-))dt + b(\mathbf{q}(t^-), \mathbf{x}(t^-))d\mathbf{w}(t) \\ &+ \int_{\mathbb{R}^n} \int_0^1 \int_{\Gamma} \theta(z, u, \mathbf{q}(t), (\mathbf{q}(t^-), \mathbf{x}(t^-)))m(\mathbf{q}(t^-), \mathbf{x}(t^-), \gamma)\mu(dz \times du \times d\gamma, dt). \end{aligned} \tag{1.14}$$

Notice that $\bar{\xi}\widehat{\lambda} \geq \widehat{\lambda}$ (and thus $\widehat{\Delta} \subseteq \Theta$), which means that the Poisson random measure μ, acting in equation (1.14), elicits arrivals more frequently than that in (1.11). In order to consider an interconnected model for the dynamics (1.11-1.14) of the SHS $\mathscr{S}_{\lambda}$ that would be equivalent to (1.8-1.9), we ought to have Poisson arrivals within p, μ with the same rates. In other words, let us find a way to "synchronize" the random events in time. Let us rescale the intensity of the random measure $p(d\gamma, dt)$ in (1.11) to be $\mathscr{U}_2(d\gamma) \times \bar{\xi}\widehat{\lambda}dt$, where $\mathscr{U}_2$ is uniformly distributed in Θ (the new mark space).

Notice that in principle using a higher arrival rate for the introduced random measures, while keeping the transition intensities unchanged, is computationally undesirable. In fact, that does not increase the likelihood of jumping from a specific point in the hybrid state space to a certain mode, or the likelihood that a jump occurs at all. More precisely, let us consider the random measure $p(d\gamma, dt)$: it assigns a mark in $\gamma \in \Theta$ at time $\tau > 0$ if there exists $n \geq 1$ such that the arrival time t_n is equal to τ and the point γ is sampled from the distribution $\mathscr{U}_2(\Gamma)$. Then, the probability in

(1.9), for $\forall j \in \mathcal{Q}, j \neq i$, can be expressed as:

$$\begin{aligned}&\mathcal{P}(\mathbf{q}(t+dt)=j|\mathbf{q}(t)=i,\mathbf{x}(t)) = p(\Delta_{ij}(\mathbf{x}(t)) \times [t,t+dt])\\ &= \textstyle\sum_{n\geq 1} \mathbf{1}_{\{t\leq t_n\leq t+dt\}}\mathbf{1}_{\Delta_{ij}(\mathbf{x}(t))} \approx \widehat{\xi}\widehat{\lambda}dt\frac{\lambda_{ij}(\mathbf{x}(t))}{\widehat{\xi}\widehat{\lambda}} = \lambda_{ij}(\mathbf{x}(t))dt,\end{aligned}$$

where $\mathbf{1}_{\{.\}}$ is the indicator function for the measurable event described in its argument. The above probability is approximately the expression that was obtained in (1.9), and is thus independent of the actual arrival rate. This oversampling effect can be interpreted as the creation of a number of "void samples," which will regardless employ some computational time to be processes and discarded. In principle, it would then be preferable to use random arrivals with the lowest possible intensities. However, notice that in the present instance it would not be possible to "undersample" the arrivals of the random measure μ, because that would affect the distribution of the resets.

To summarize, the global dynamics of $\mathscr{S}_\lambda$ will be described, at any time $t \geq 0$, by a process $\mathbf{s}(t) = (\mathbf{q}(t), \mathbf{x}(t))$ which, given an initial condition $\mathbf{s}(0) = (\mathbf{q}(0), \mathbf{x}(0)) = (q_0, x_0) \in \mathcal{S}$, is the solution of the following set of relations:

$$\left\{\begin{array}{l}\mathbf{q}(t) = q_0 + \int_0^t \int_\Gamma l(\mathbf{q}(s^-), \mathbf{x}(s^-), \gamma)p(d\gamma, ds),\\ \mathbf{x}(t) = x_0 + \int_0^t a(\mathbf{q}(s^-), \mathbf{x}(s^-))ds + \int_0^t b(\mathbf{q}(s^-), \mathbf{x}(s^-))d\mathbf{w}(s)\\ + \int_0^t \int_{\mathbb{R}^n} \int_0^1 \int_\Gamma \theta(z, u, q(s), (\mathbf{q}(s^-), \mathbf{x}(s^-)))m(\mathbf{q}(s^-), \mathbf{x}(s^-), \gamma)\mu(dz \times du \times d\gamma, ds).\end{array}\right. \tag{1.15}$$

Notice the dependence, along time, of the second dynamical relation on the value of the first one. This dependence follows from the semantics of the SHS model $\mathscr{S}_\lambda$, and will be reflected in some specific terms of the time-discretized version of the dynamics in (1.15).

Remark 9. *The formulation in (1.15) is rather general and concise. It provides an advancement from that in [Ghosh and Bagchi, 2004; Krystul and Bagchi, 2004], because the resets in those works are described by deterministic functions, while here they explicitly depend on stochastic kernels, which is translated in the necessity to introduce the function θ.*

The model can be generalized to the case where the jumps in the continuous and

the discrete parts are not synchronized, as further pursued in [Ghosh and Bagchi, 2004]. Also, the assumption that imposes the same domain dimension can be relaxed, provided some heavier notations are traded off. In particular, it would be necessary to introduce intervals of the real line (like $\Delta_{ij}(x)$) that are domain-characteristic. □

The following result nicely connects with Proposition 1. The proposition is based on [Blom, 2003], where a proof of the statement can be found.

Proposition 5 (Existence and Uniqueness of Solutions)**.** *Assume the conditions 1.1, 1.3, 1.6, 1.7, 1.8 are valid. Let 1.5 hold in terms of the standard Wiener process, and assume further independence with respect to the introduced random measures p, μ. Then the system in (1.15) has a unique global solution.* □

Alternative Formulation. Consider the function l defined above. It has been built by sampling uniformly over an interval built by summing the contributions of the single transition intensities. It would be desirable to transpose this approach to the choice of the reset function, modulo the understanding that they do not affect the actual jumping conditions, but are instead associated with any jump by influencing the reset of the hybrid process.

Recall the introduction of functions $\tilde{f}$ and, by the definition of $\tilde{\xi}$, of functions g. Now, modify the definition of $\breve{f} : \mathbb{R}^n \to \mathbb{R}^+$ as $\breve{f}(z) = \sum_{q' \in \mathcal{Q}} \tilde{f}(q', z)$ and of the constant $\breve{\xi} = \sum_{q' \in \mathcal{Q}} \tilde{\xi}(q')$. Finally, $\check{f} : \mathbb{R}^n \to \mathbb{R}^+$ is defined as $\check{f}(z) = \frac{\breve{f}(z)}{\breve{\xi}}$.

Introduce the following function $\theta : \mathbb{R}^n \times \mathbb{R}^+ \times \mathcal{Q} \times \mathcal{S} \to \{0, 1\}$ as:

$$\theta(z, u, q', (q, x)) = \begin{cases} i - j, & \text{if } \dfrac{\sum\limits_{j<q', j \neq q} R^{\Lambda}(j, z, (q, x))}{\breve{\xi}\breve{f}(z)} \leq u < \dfrac{\sum\limits_{j \leq q', j \neq q} R^{\Lambda}(j, z, (q, x))}{\breve{\xi}\breve{f}(z)}, \\ 0, & \text{else.} \end{cases}$$

By introduction of the Kronecker delta function over the integers, define

$$\delta(l(q, x, \gamma) + \theta(z, u, q', q, x)) = \begin{cases} 1, & \text{if } \gamma \in \Delta_{qq'}(x) \wedge \\ & \dfrac{\sum\limits_{j<q', j \neq q} R^{\Lambda}(j, z, (q, x))}{\breve{\xi}\breve{f}(z)} \leq u < \dfrac{\sum\limits_{j \leq q', j \neq q} R^{\Lambda}(j, z, (q, x))}{\breve{\xi}\breve{f}(z)}, \\ 0, & \text{else.} \end{cases}$$

Introduce now two random measures: $\mu(dz, du, d\gamma, dt)$, with intensity $\check{f}(z) \times \mathscr{U}_1(du) \times \mathscr{U}_2(d\gamma) \times \check{\xi}\widehat{\lambda}dt$; $p(d\gamma, dt)$, with intensity $U_2(d\gamma) \times \check{\xi}\widehat{\lambda}dt$, which will be oversampled to achieve synchronization with the first. Then, the dynamics can be expressed as follows:

$$\begin{cases} \mathbf{q}(t) = q_0 + \int_0^t \int_\Gamma l(\mathbf{q}(s^-), \mathbf{x}(s^-), \gamma) p(d\gamma, ds), \\ \mathbf{x}(t) = x_0 + \int_0^t a(\mathbf{q}(s^-), \mathbf{x}(s^-)) ds + \int_0^t b(\mathbf{q}(s^-), \mathbf{x}(s^-)) d\mathbf{w}(s) \\ + \sum_{q' \in \mathcal{Q}} \int_0^t \int_{\mathbb{R}^n} \int_0^1 \int_\Gamma \delta(l(\mathbf{q}(s^-), \mathbf{x}(s^-), \gamma) + \theta(z, u, q', (\mathbf{q}(s^-), \mathbf{x}(s^-))) \mu(dz \times du \times d\gamma, ds). \end{cases}$$

Notice that, because $\check{\xi} \geq \bar{\xi}$, the expression in (1.15) is computationally more attractive than that just above. In the following, let us work with the model in (1.15).

The introduced structure for the dynamics will be used to define a time discretization scheme in section 1.4.2.

1.4.2 State-Dependent Thinning Procedure

The random hybrid dynamics of the SHS $\mathscr{S}_\lambda$ have been succinctly expressed via the system of equations in (1.15). The convenience of this form is not only aesthetic. It is also important to stress that the random measures p, μ that drive the model are Poisson point processes. In other words, the arrivals occur with an intensity that is independent of the hybrid state space. This holds for both the continuous and the discrete dynamics. Section 1.4.1 has shown a rather general procedure to turn a non-homogeneous arrival and jump processes into one of Poisson kind. This independence enables one to employ an array of time sampling methods which, provided some regularity conditions on the elements of the model, allow one to prove some sort of convergence between the discrete-time sampled process and the original continuous-time process. Intuitively, sampling a Poisson arrival process does not require the knowledge, along time, of the stochastic process to be sampled. Let us focus, to begin with, on the control-independent case.

In order to obtain a discretization in time of the above dynamics, let us start by introducing a sequence of discrete instants $\{\tau_i\}, i \in \mathcal{I} = \{0, 1, \ldots, N\}$, over a finite time horizon $[0, T]$. This sequence is made up of the deterministic grid $\{\tilde{\tau}_i\}, i =$

$0, 1, \ldots, \tilde{N} = \left\lfloor \frac{T}{\delta t} \right\rfloor \leq N$, obtained by uniformly sampling the interval $[0, T]$ with time step $\delta t > 0$, and of the random switching instants $\{\bar{\tau}_i\}, i = 1, \ldots, (N - \tilde{N})$ over $[0, T]$. It then holds that $\{\tau_i\} = \{\widehat{\tau}_i\} \bigcup \{\bar{\tau}_i\}$. It is important to notice that the random Poisson times can be computed independently of the knowledge of the actual trajectory, because of the state-independent arrival rates that make up the Poisson random measures in (1.15).

The values of the discretized process at the time instants $\{\tau_i\}$ are obtained via a (first-order) Euler scheme applied to $\mathbf{s}(t), t \in [0, T]$. As a reminder, notice that the càdlàg process $\mathbf{q}(t)$ is piecewise constant, which implies that $\mathbf{q}(\tau_{i+1}^-) = \mathbf{q}(\tau_i)$.

Given an initial condition $\mathbf{s}(0) = (\mathbf{q}(0), \mathbf{x}(0)) = (q_0, x_0)$ and assuming by default $\tau_0 = 0$, the following iterative scheme is obtained on the time instants set $\{\tau_i\}, i \in \mathcal{I}$, which yields the discrete-time process $\widehat{\mathbf{s}}(\tau_i) = (\widehat{\mathbf{q}}(\tau_i), \widehat{\mathbf{x}}(\tau_i)), i \in \mathcal{I}$:

$$\widehat{\mathbf{x}}(\tau_{i+1}^-) = \widehat{\mathbf{x}}(\tau_i) + a(\widehat{\mathbf{q}}(\tau_i), \widehat{\mathbf{x}}(\tau_i))(\tau_{i+1} - \tau_i) + b(\widehat{\mathbf{q}}(\tau_i), \widehat{\mathbf{x}}(\tau_i))(\widehat{\mathbf{w}}(\tau_{i+1}) - \widehat{\mathbf{w}}(\tau_i)) \tag{1.16}$$

$$\widehat{\mathbf{q}}(\tau_{i+1}) = \widehat{\mathbf{q}}(\tau_i) + \int_\Gamma l(\widehat{\mathbf{q}}(\tau_i), \widehat{\mathbf{x}}(\tau_{i+1}^-), z) p(dz, \tau_{i+1}); \tag{1.17}$$

$$\widehat{\mathbf{x}}(\tau_{i+1}) = \widehat{\mathbf{x}}(\tau_{i+1}^-) + \int_{\mathbb{R}^n} \int_0^1 \int_\Gamma \theta(z, u, \widehat{\mathbf{q}}(\tau_{i+1}), (\widehat{\mathbf{q}}(\tau_{i+1}^-), \widehat{\mathbf{x}}(\tau_{i+1}^-))) \cdot \cdot m(\widehat{\mathbf{q}}(\tau_{i+1}^-), \widehat{\mathbf{x}}(\tau_{i+1}^-), \gamma) \mu(dz, du, d\gamma, \tau_{i+1}). \tag{1.18}$$

In these formulas the random measures are evaluated at the jump epochs. In order to interpret and understand the sequence of the discretization steps above, it is useful to refer to the original semantics of the hybrid trajectory of a SHS, described in Definition 6. To begin with, the dynamics of the continuous component are updated according to the contributions of the drift and diffusion terms in equation (1.16). Equation (1.17), which depends on (1.16), updates the discrete value of the mode according to the Poisson random measure $p(dz, dt)$. Finally, the relation in (1.18) updates the value of the continuous dynamics only at the jump time, according to the other Poisson random measure $\mu(dz, du, d\gamma, dt)$. Notice that equation (1.18) depends on the new value assumed in (1.17).

1.4.3 A Discretization Scheme with Convergence

As discussed while introducing the concept of process generator and presenting the Dynkin formula (equation (1.5) in section 1.2.2), it is generally interesting to compute expectations of functions of the stochastic processes under study. This connects to the notion of weak convergence, introduced in Definitions 10 and 11, which describes the limiting behavior of the expectation of a sequence of functions of the process under study. Here, given the original stochastic process $\mathbf{s}(t), t \geq 0$, solution of the SHS $\mathscr{S}_\lambda$, and the time-discretized process $\widehat{\mathbf{s}}(\tau_i), i \in \mathcal{I}$, obtained by the application of a discretization scheme as obtained in (1.16-1.17-1.18) with discretization step δt, given any continuous and bounded function $f : \mathcal{S} \to \mathbb{R}$, it would be desirable to show that

$$|E[f(\mathbf{s}(\tau_i))] - E[f(\widehat{\mathbf{s}}(\tau_i))]| \leq \beta(\delta t)^\alpha,$$

where $0 < \alpha, \beta < \infty$. In particular, if the above bound holds, then α is known as the *order* of convergence.

The convergence result of the above thinning procedure that will be presented now for a time discretization procedure based on the Euler scheme [Milstein, 1994] is drawn from the works [Glasserman and Merener, 2003; Glasserman and Merener, 2004]. These studies are in turn based on results from [Kloeden and Platen, 1992]. The two references [Glasserman and Merener, 2003; Glasserman and Merener, 2004] are of interest here because their results have been derived cognizant of the fact that some functions appearing in the dynamics (here, the terms $\theta(\cdot), l(\cdot), m(\cdot)$) are inescapably discontinuous. For reference, the Euler scheme here introduced coincides with the known Euler-Maruyuama scheme for diffusions, which will be further exploited in section 2.3. Let us introduce the indicator function $\mathbf{1}_{\mathcal{S}}(s) = 1, \forall s \in \mathcal{S}$, and for simplicity rename $c(z, u, \gamma, q', s) \doteq \theta(z, u, q', s)m(s, \gamma)\mathbf{1}_{\mathcal{S}}(s)$.To leverage some results for martingales, introduce the quantity

$$\tilde{a}(s) = a(s) + \int_{\mathbb{R}^n} c(\cdot, z) f(s) \widehat{\xi\lambda} dz,$$

where f is the intensity function introduced for equation (1.12), given the validity of Assumption 6. Consider a bounded function $\psi : \mathbb{R}^n \to \mathbb{R}$. Introduce the two

functions, for $x \in \mathbb{R}^n$:

$$\begin{aligned} \varphi_c(x) &= \int_{\mathbb{R}^n} \psi(x + c(\cdot, z, x)) f(z) dz; \\ \bar{\varphi}_c(x) &= \int_{\mathbb{R}^n} (x + c(\cdot, z, x)) f(z) dz. \end{aligned}$$

The following holds:

Theorem 4 (Weak Convergence of the Euler Scheme)**.** *Consider a real-valued function defined on the hybrid state space, $g : \mathcal{S} \to \mathbb{R}$, and assume $g \in \bar{C}_b^4(\mathcal{S})$. Let $\mathbf{s}(t), t \in [0, N]$ be the solution of (1.15), and $\widehat{\mathbf{s}}(\tau_i), i \in \mathcal{I}$, be the solution of the first-order Euler scheme in (1.16-1.17-1.18). Assume that*

- *$\bar{\varphi}_c(x) \in \bar{C}_b^4(\mathbb{R}^n)$;*
- *If, for some $0 < \Psi < \infty$, $\psi \in \bar{C}_b^4(\mathbb{R}^n)$ with bound Ψ, then there exists a constant $0 < K < \infty$, such that $\varphi_c \in \bar{C}_b^4(\mathbb{R}^n)$ with bound $K\Psi$;*
- *For any $q \in \mathcal{Q}$, the functions $\tilde{a}(q, x), b(q, x) \in \bar{C}_b^4(\mathbb{R}^n)$;*
- *There exists a constant $0 < H < \infty$ such that, for any $q \in \mathcal{Q}, \|\tilde{a}(q, y)\| \leq H(1 + \|y\|)$, and $\|b(q, y)\| \leq H(1 + \|y\|), \forall y \in \mathbb{R}^n$.*

Then, $\mathbf{s}(\cdot)$ has weak convergence of order one to $\widehat{\mathbf{s}}(\cdot)$. In other words, for any such g, there exists a quantity $0 < \beta < \infty$ such that, for a small enough integration step δt,

$$|E[g(\mathbf{s}(\tau_i))] - E[g(\widehat{\mathbf{s}}(\tau_i))]| \leq \beta\, \delta t, \quad i \in \mathcal{I}. \qquad \square$$

Proof. The proof is an extension of that in [Glasserman and Merener, 2004, Theorem 6.1], which is also discussed in [Glasserman and Merener, 2003] and based on results in [Mikulevicius and Platen, 1988].

The main idea comes from realizing that, although some elements in the dynamics are possibly discontinuous $(\theta(\cdot), l(\cdot), m(\cdot))$, their discontinuities will possibly happen only at the switching times $\tau_i, i \in \mathcal{I}$. More precisely, for a small enough integration step (that is, for a small enough interval $[\tau_i, \tau_{i+1}), i \in \mathcal{I}$), the discontinuous thinning functions $c(\cdot, x), l(\cdot, x), x \in \mathbb{R}^n$ will not change their constant value

within the integration interval $[\tau_i, \tau_{i+1})$. That is, the discontinuities of the functions $c(\cdot, x), l(\cdot, x), x \in \mathbb{R}^n$ may happen, for small enough time intervals, exclusively at the jumping instants, for any of the values of the samples of the two uniform distributions introduced in equations (1.11) and (1.13).

This argument allows one to prove the first continuity requirement in the statement for the functions into play. The third condition requires a strengthening of Assumption 1. In particular, the continuity of the function $\tilde{a}$ requires one to proceed as done for the first point.

The second condition is proven along similar lines and much like in [Glasserman and Merener, 2004, Theorem 6.1], and again hinges on arguments that are similar as above.

Finally, the last condition follows from the Lipschitz assumption on the drift and diffusion terms, as explained in equation (1.2), Assumption 1.1.

These observations allow one to relate the conditions of the statement to [Mikulevicius and Platen, 1988, Theorem 3.3], which is then leveraged to actually prove the claim. □

Remark 10. *The result has been obtained with a continuity requirement on the functions g. Let us stress that, while the result of the above statement does not strictly imply the weak convergence of the associated processes, it is of interest to the current study because we shall be working with expectations of the trajectories of the solutions to the SHS. Notice also that it is possible to apply discretization schemes of higher order [Milstein, 1994], as suggested in [Glasserman and Merener, 2004]. Their weak convergence can be as well proven, albeit at the expense of introducing stronger regularity assumptions.* □

Chapter 2

Probabilistic Reachability and Safety

2.1 The Concept of Reachability in Systems and Control Theory

Reachability is an important topic in classical control theory. In general terms, a reachability problem consists of evaluating whether the state of a system will reach a certain "target" set during some time horizon, starting from a given set of initial conditions and possibly subject to a control input (see Figure 2.1).

This problem is of interest, for instance, in those safety problems where the system should be kept outside of an unsafe region of the state space and the control input can be chosen so as to avoid this unsafe region (darker, blue region in Figure 2.2). In a deterministic setting, reachability is a yes/no problem, where one evaluates if starting from a given set of initial states the system will reach a certain set or not. In a stochastic setting, the different trajectories originating from each initial state have a different likelihood and one can then evaluate what is the probability that the system will reach the assigned set starting from a certain initial distribution over the set of initial states. In safety problems where the evolution of the system can be influenced by some control input, one should select it appropriately so as to minimize the probability that the state of the system will enter the unsafe set.

Much investigation has been done on reachability analysis for system verification, where the problem consists in verifying if some designed system satisfies certain

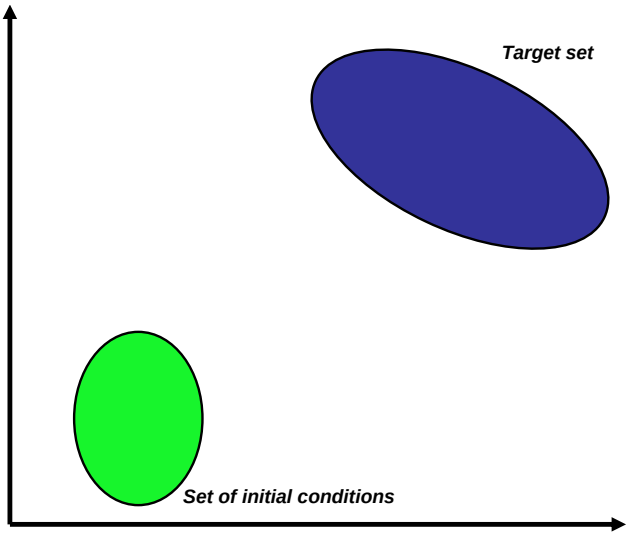

Figure 2.1: The reachability problem, a pictorial representation of the entities into play.

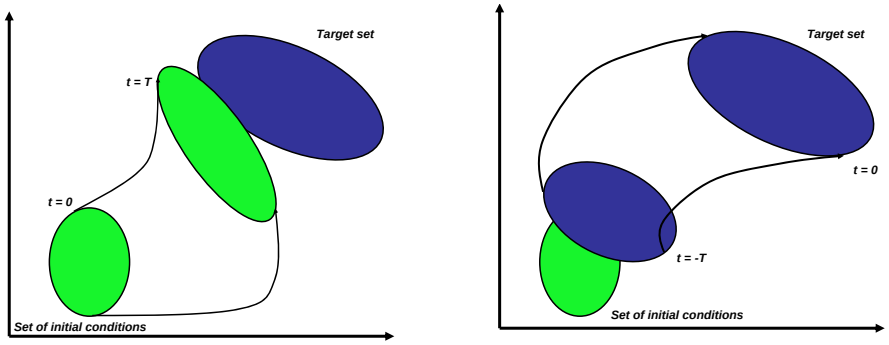

Figure 2.2: The reachability problem, two intuitive interpretations of its computational verification.

reachability specifications encoding a correct/desired behavior.

In the case of deterministic systems, 'model checking' is the most commonly adopted technique for system verification, where reachability specifications are verified by constructing reachable sets based on a model of the system (either propagating "forward," or "backward" the flow of the system—see Figure 2.2). This is possible only for simple dynamics [Anai and V.Weispfenning, 2001; Lafferriere *et al.*, 2000; Lafferriere *et al.*, 2001]. In the hybrid systems case, set representation and propagation by continuous flow is generally difficult, and termination of the algorithm for reachable set computation is not guaranteed since the state space is not finite, [Tomlin *et al.*, 2003]. Decidability results have been proven to hold only for certain classes of hybrid systems by using discrete abstraction consisting in building a finite automaton that is 'equivalent' (bisimilar) to the original hybrid system for the purpose of verification [Alur *et al.*, 2000].

In the case of complex dynamics, some approximation methods are needed for reachability computations. Two main approaches have been introduced to this purpose: seeking for an abstraction of the system that would yield a simpler model for solving the original reachability problem, or adopting an approximation of sets that can be easily represented and propagated through the system dynamics. In the first approach, an approximate simulation relation is introduced to obtain an abstraction of the original system [Girard *et al.*, 2006], while a quantifier elimination decision procedure and simple theorem proving is introduced in [Tiwari and Khanna, 2002; Tiwari and Khanna, 2004] to construct a series of successively finer qualitative abstractions of the hybrid automaton under study. In the second approach, over-approximations by ellipsoids [Kurzhanski and Varaiya, 2000; Kurzhanski and Varaiya, 2002], polyhedra [Asarin *et al.*, 2000; Asarin *et al.*, 2003], zonotopes [Girard, 2005], oriented rectangular polytopes [Stursberg and Krogh, 2003; Yazarel and Pappas, 2004], general polygones [Chutinan and Krogh, 1998; Han and Krogh, 2006; Kerrigan *et al.*, 2006] were proposed, or, alternatively, asymptotic approximations of reachable sets that converge to the true reachable sets as some accuracy parameter tends to zero. Level set methods [Mitchell *et al.*, 2001; Mitchell and Tomlin, 2000] and gridding [Belta *et al.*, 2004] techniques belong to this latter category.

A connection of reachability (and related concepts, such as safety or viability)

with optimal control for deterministic problems has been pointed out in [Hedlund and Rantzer, 2002; Lygeros, 2004b]. The connection between reachability, safety and dynamic games for deterministic hybrid systems has been stressed in [Mitchell *et al.*, 2005; Lygeros *et al.*, 1999; Tomlin *et al.*, 1998a], where it is mostly applied to air traffic management problems.

Reachability for stochastic hybrid systems has been a very recent focus of research. Most of the approaches consider the problem of reachability analysis for continuous time stochastic hybrid systems (CTSHS) without a control input. The theory of CTSHS, progressively developed since the early contributions in [Davis, 1993; Ghosh *et al.*, 1997; Hu *et al.*, 2000], is used in [Bujorianu and Lygeros, 2003] to address theoretical issues regarding the measurability of the reachability events. In [Bujorianu, 2004], the theory of Dirichlet forms associated with a right-Markov process is employed in studying a probabilistic reachability problem, and upper bounds for the reach set probabilities are derived.

The contributions in [Hu *et al.*, 2003; Hu *et al.*, 2005; Prandini and Hu, 2006a; Prandini and Hu, 2006b] address the reachability problem using a Markov chain approximation, [Kushner and Dupuis, 2001], to compute the probability of entering some assigned set, and apply the concept to air traffic control studies. Probabilities rather than sets are propagated through the approximating Markov chain transition kernel. In the same spirit, model checkers for verifying probabilistic reachability specifications of Markov chains have been developed, [Katoen, 2006]. From a different perspective, in [Prajna *et al.*, 2004] certain functions of the state of the system known as barrier certificates are used to compute an upper bound on the probability of reaching a set. These barrier functions are synthesized by a sum-of-squares SOS method [Yazarel *et al.*, 2004]. The approach in [Digailova and Kurzhanski, 2005] is unique in introducing a 'mean-square' definition for the concept of reachability. In [Mitchell and Templeton, 2005] the control case is considered in a rather general game theoretical framework, and a reachability problem is introduced as the solution of a Hamilton-Jacobi-Isaacs partial differential equation. [Lygeros and Watkins, 2003; Prandini *et al.*, 2000] compute the reachability probability using randomized algorithms, motivated by air traffic control applications.

In this study let us adopt a discrete time point of view in order to gain a deeper

understanding of the theoretical and computational aspects associated with reachability and safety problems for stochastic hybrid systems, while avoiding technical measurability issues that arise in the continuous time case. In particular, we develop a methodology to compute and maximize the probability of maintaining the state of the system within a certain 'safe' region for a class of discrete time stochastic hybrid system (DTSHS) whose dynamics can be influenced by a control input. Unlike previous approaches, the safe set can be time-varying, which allows us to generalize the approach towards problems of regulation and stabilization, [Bertsekas, 1972; Picasso and Bicchi, 2005], by appropriately reinterpreting them as safety problems.

The proposed methodology is based on formulating the reachability problem as a stochastic optimal control problem. Based on the expression of the probability that the state of the controlled system will evolve within the safe region as a multiplicative cost, dynamic programming (DP) can be used to compute the Markov policy maximizing the cost, and also the maximally safe sets corresponding to different safety levels. These are the set of initial conditions for the system, such that there exists a Markov policy capable of maintaining the state of the system within the safe set with a probability greater than a prescribed safety level (see [Tomlin *et al.*, 1998a; Balluchi *et al.*, 2000] for a similar notion in the deterministic case).

Adopting a dual perspective, where the objective is that of minimizing the probability that the system will exit the safe set, let us again formulate the reachability problem as a stochastic optimal control problem, but this time with a cost that is the maximum of a function of the state over the time horizon. DP is shown to be still effective in this case to determine probabilistic maximal safe sets for Markov policies. In fact, the value functions for the multiplicative cost and the max cost can be properly put in relation, thus formalizing the intuition that the two viewpoints for reachability are complementary to each other.

Connecting with the simple examples introduced in chapter 1, in the second part of this section a room temperature regulation problem shall be used as running example to illustrate the DTSHS model formalism and the approach to reachability. This case study is inspired by one of the benchmark problems proposed in [Fehnker and Ivančić, 2004].

2.1.1 Discrete-Time Controlled Stochastic Hybrid Systems

Let us define a DTSHS as the discrete time counterpart of the general continuous time model described in section 1.2.1, which extends in expressiveness previous continuous time models [Davis, 1993; Ghosh *et al.*, 1997; Hu *et al.*, 2000]. More precisely, the model about to be introduced can be thought of as being obtained from the GSHS model by a time-discretization procedure as that in section 1.4, with the addition of controls. However, for the sake of clarity, let us formally describe this hybrid model and the properties which are of of nearest interest to the current study. In particular, let us stress that, unlike the GSHS models introduced in the preceding part, the following model shall embed the knowledge of the transition *probabilities* directly in its definition, rather than passing through the use of transition intensities. This is motivated by the presence of the discrete time — notice that the time-discretization procedure reinterpreted the presence of transition intensities as rates of arrival associated with a proper point process.

The state of a DTSHS is characterized by a discrete and a continuous component. The discrete state component takes on values in a countable set of modes $\mathcal{Q}$. The continuous state space in each mode $q \in \mathcal{Q}$ is given by the Euclidean space $\mathbb{R}^{n(q)}$, whose dimension $n(q)$ is determined by the map $n : \mathcal{Q} \to \mathbb{N}$. Thus the hybrid state space is $\mathcal{S} := \bigcup_{q\in\mathcal{Q}}\{q\} \times \mathbb{R}^{n(q)}$. Let $\mathcal{B}(\mathcal{S})$ be the σ-field generated by the subsets of $\mathcal{S}$ of the form $\bigcup_q\{q\} \times A_q$, where A_q is a Borel set in $\mathbb{R}^{n(q)}$. The fact that $\mathcal{S}$ can be endowed with a metric that is equivalent to the usual Euclidean metric when restricted to each domain $\mathbb{R}^{n(q)}$, [Davis, 1993], shows that $(\mathcal{S}, \mathcal{B}(\mathcal{S}))$ is a Borel space, i.e. homeomorphic to a Borel subset of a complete separable metric space.

The continuous state of a DTSHS evolves according to a probabilistic law that depends on the actual operating mode. A discrete transition from the current operating mode to a different one may occur during the continuous state evolution, again according to some probabilistic law. This will in turn cause a modification of the probabilistic law governing the continuous state dynamics. A control input can affect the discrete and continuous evolution of the system. After a discrete transition has occurred, the continuous state component is subject to a probabilistic reset that is also influenced by some control input. Let us distinguish this latter input from the

former one, naming them respectively *reset* and *transition input*.

Definition 13 (Discrete-Time, controlled Stochastic Hybrid System)**.** *A discrete time stochastic hybrid system (DTSHS) is a tuple* $\mathscr{S} = (\mathcal{Q}, n, \mathcal{U}, \Sigma, T_x, T_q, R)$*, where*

- $\mathcal{Q} := \{q_1, q_2, \ldots, q_m\}$*, for some* $m \in \mathbb{N}$*, represents the discrete state space;*
- $n : \mathcal{Q} \to \mathbb{N}$ *assigns to each discrete state value* $q \in \mathcal{Q}$ *the dimension of the continuous state space* $\mathbb{R}^{n(q)}$*. The hybrid state space is then given by* $\mathcal{S} := \bigcup_{q \in \mathcal{Q}} \{q\} \times \mathbb{R}^{n(q)}$*;*
- $\mathcal{U}$ *is a compact Borel space representing the transition control space;*
- Σ *is a compact Borel space representing the reset control space;*
- $T_x : \mathcal{B}(\mathbb{R}^{n(\cdot)}) \times \mathcal{S} \times \mathcal{U} \to [0,1]$ *is a Borel-measurable stochastic kernel on* $\mathbb{R}^{n(\cdot)}$ *given* $\mathcal{S} \times \mathcal{U}$*, which assigns to each* $s = (q, x) \in \mathcal{S}$ *and* $u \in \mathcal{U}$ *a probability measure on the Borel space* $(\mathbb{R}^{n(q)}, \mathcal{B}(\mathbb{R}^{n(q)}))$*:* $T_x(\cdot|s, u)$*;*
- $T_q : \mathcal{Q} \times \mathcal{S} \times \mathcal{U} \to [0,1]$ *is a discrete stochastic kernel on* $\mathcal{Q}$ *given* $\mathcal{S} \times \mathcal{U}$*, which assigns to each* $s \in \mathcal{S}$ *and* $u \in \mathcal{U}$*, a probability distribution over* $\mathcal{Q}$*:* $T_q(\cdot|s, u)$*;*
- $R : \mathcal{B}(\mathbb{R}^{n(\cdot)}) \times \mathcal{S} \times \Sigma \times \mathcal{Q} \to [0,1]$ *is a Borel-measurable stochastic kernel on* $\mathbb{R}^{n(\cdot)}$ *given* $\mathcal{S} \times \Sigma \times \mathcal{Q}$*, that assigns to each* $s \in \mathcal{S}$*,* $\sigma \in \Sigma$*, and* $q' \in \mathcal{Q}$*, a probability measure on the Borel space* $(\mathbb{R}^{n(q')}, \mathcal{B}(\mathbb{R}^{n(q')}))$*:* $R(\cdot|s, \sigma, q')$*.* □

To describe the semantics of $\mathscr{S}$, an initial condition $s_0 \in \mathcal{S}$ needs to be specified, as well as how the reset and transition inputs are chosen.

The system initialization can be specified through some probability measure $\pi \in \mathscr{P}(\mathcal{S}), \pi : \mathcal{B}(\mathcal{S}) \to [0,1]$ on the Borel space $(\mathcal{S}, \mathcal{B}(\mathcal{S}))$. When the initial state of the system is $s \in \mathcal{S}$, then, the probability measure π is concentrated at $\{s\}$. As for the choice of the reset and transition inputs, we need to specify which is the rule to determine their values at every time step during the DTSHS evolution (control policy). Here, let us consider a DTSHS evolving over a finite time horizon $[0, N]$, with inputs chosen according to a Markov policy. If the values for the control inputs at each time $k \in [0, N)$ are determined based on the values taken by the past inputs and the state up to the current time k, then the policy is said to be a feedback policy.

Definition 14 (Feedback policy)**.** *Let $\mathscr{S} = (\mathcal{Q}, n, \mathcal{U}, \Sigma, T_x, T_q, R)$ be a DTSHS. A feedback policy μ for $\mathscr{S}$ is a sequence $\mu^f = (\mu_0^f, \mu_1^f, \ldots, \mu_{N-1}^f)$ of universally measurable maps $\mu_k^f : \mathcal{S} \times (\mathcal{S} \times \mathcal{U} \times \Sigma)^k \to \mathcal{U} \times \Sigma$, $k = 0, 1, \ldots, N-1$.* □

Let us denote the set of feedback policies as $\mathcal{M}$.

If the values for the control inputs are determined only based on the value taken by the state at the current time step, i.e., $(u_k, \sigma_k) = \mu_k(s_k)$, then the policy is said to be a Markov policy.

Definition 15 (Markov Policy)**.** *Consider a DTSHS $\mathscr{S} = (\mathcal{Q}, n, \mathcal{U}, \Sigma, T_x, T_q, R)$. A Markov policy μ for $\mathscr{S}$ is a sequence $\mu = (\mu_0, \mu_1, \ldots, \mu_{N-1})$ of universally measurable maps*

$$\mu_k : \mathcal{S} \to \mathcal{U} \times \Sigma, \quad k = 0, 1, \ldots, N-1,$$

from the hybrid state space $\mathcal{S} = \bigcup_{q \in \mathcal{Q}} \{q\} \times \mathbb{R}^{n(q)}$ to the control input space $\mathcal{U} \times \Sigma$. □

Let us denote the set of Markov policies as $\mathcal{M}_m$. Clearly $\mathcal{M}_m \subseteq \mathcal{M}$. In this work, let us focus on Markov policies.

Let us recall that a function $\mu_k : \mathcal{S} \to \mathcal{U} \times \Sigma$ is universally measurable if the inverse image of every Borel set is measurable with respect to every complete probability measure on $\mathcal{S}$ that measures all Borel subsets of $\mathcal{S}$. This measurability condition is weaker than the Borel measurability condition and is needed to assess properties which hold uniformly in the initial condition $s_0 \in \pi$, [Bertsekas and Shreve, 1996]. In practice, this assumption is needed for showing the existence of a solution of an optimization problem that will be set up in the following.

The semantics of a DTSHS can be algorithmically defined through the notion of *execution*. In the rest of this work, we shall use boldface to denote random variables and normal typeset to denote sample values.

Definition 16 (DTSHS Execution)**.** *Consider a DTSHS $\mathscr{S} = (\mathcal{Q}, n, \mathcal{U}, \Sigma, T_x, T_q, R)$ and a time horizon $[0, N]$. A stochastic process $\{\mathbf{s}(k) = (\mathbf{q}(k), \mathbf{x}(k)),\ k \in [0, N]\}$ with values in $\mathcal{S} = \bigcup_{q \in \mathcal{Q}} \{q\} \times \mathbb{R}^{n(q)}$ is an execution of $\mathscr{H}$ associated with a Markov policy $\mu \in \mathcal{M}_m$ and an initial condition $s_0 = (q_0, x_0) \in \mathcal{S}$ if its sample paths are obtained according to the DTSHS algorithm:*

Algorithm 3.
set $\mathbf{q}(0) = q_0$, $\mathbf{x}(0) = x_0$, *and* $k = 0$;

while $k < N$ *do*

set $(u_k, \sigma_k) = \mu_k((q_k, x_k))$;

extract from $\mathcal{Q}$ *a value* q_{k+1} *for* $\mathbf{q}(k+1)$ *according to* $T_q(\cdot\,|(q_k, x_k), u_k)$;

if $q_{k+1} = q_k$, *then*

extract from $\mathbb{R}^{n(q_{k+1})}$ *a value* x_{k+1} *for* $\mathbf{x}(k+1)$ *according to* $T_x(\cdot\,|(q_k, x_k), u_k)$;

else

extract from $\mathbb{R}^{n(q_{k+1})}$ *a value* x_{k+1} *for* $\mathbf{x}(k+1)$ *according to* $R(\cdot\,|(q_k, x_k), \sigma_k, q_{k+1})$;

$k \to k+1$;

end.

□

By appropriately defining the discrete transition kernel T_q, it is possible to model the *spontaneous jumps* that *may* occur during the continuous state evolution, as well as the *forced jumps* that *must* occur when the continuous state exits some prescribed domain.

As for the spontaneous transitions, if a discrete transition from q to $q' \neq q$ is enabled at $(q, x) \in \mathcal{S}$ by the control input $u \in \mathcal{U}$, then this can be encoded by the condition $T_q(q'|(q, x), u) > 0$.

As for the forced transitions, the *invariant set* $Inv(q)$ associated with mode $q \in \mathcal{Q}$, namely the set of all the admissible values for the continuous state within q, can be expressed in terms of T_q by forcing $T_q(q|(q, x), u)$ to be equal to zero for all the continuous state values $x \in \mathbb{R}^{n(q)}$ outside $Inv(q)$, irrespectively of the value of the control input $u \in \mathcal{U}$. Thus $Inv(q) := \mathbb{R}^{n(q)} \setminus \{x \in \mathbb{R}^{n(q)} : T_q(q|(q, x), u) = 0, \forall u \in \mathcal{U}\}$, and as soon as $x \notin Inv(q)$ while the system evolves in mode q, a jump from q to some $q' \neq q$ is forced. Unlike the continuous time model in [Bujorianu and Lygeros, 2004b], spatial guards here are implicitly defined through the map T_q. This approach is closely related to the idea of substituting the spatial guards with transition intensities (or "barriers"), but subsumes this procedure by directly defining the stochastic kernels

associated with the transition probabilities, which otherwise would have to be derived by proper integration of the transition intensities. In a discrete-time framework, the definition of the transition probabilities is certainly a more convenient choice.

Introduce the stochastic kernel $\tau_x : \mathcal{B}(\mathbb{R}^{n(\cdot)}) \times \mathcal{S} \times \mathcal{U} \times \Sigma \times \mathcal{Q} \to [0,1]$ on $\mathbb{R}^{n(\cdot)}$ given $\mathcal{S} \times \mathcal{U} \times \Sigma \times \mathcal{Q}$:

$$\tau_x(\cdot\,|(q,x),u,\sigma,q') = \begin{cases} T_x(\cdot|(q,x),u), & \text{if } q' = q \\ R(\cdot|(q,x),\sigma,q'), & \text{if } q' \neq q, \end{cases}$$

which assigns to each $s = (q,x) \in \mathcal{S}$, $u \in \mathcal{U}$, $\sigma \in \Sigma$ and $q' \in \mathcal{Q}$ a probability measure on the Borel space $(\mathbb{R}^{n(q')}, \mathcal{B}(\mathbb{R}^{n(q')}))$. τ_x is used in the DTSHS algorithm to randomly select a value for the continuous state at time $k+1$, given the values taken by the hybrid state and the control input at time k, and that of the discrete state at time $k+1$.

Based on τ_x let us introduce the Borel-measurable stochastic kernel $T_s : \mathcal{B}(\mathcal{S}) \times \mathcal{S} \times \mathcal{U} \times \Sigma \to [0,1]$ on $\mathcal{S}$ given $\mathcal{S} \times \mathcal{U} \times \Sigma$:

$$T_s((\cdot,q)\,|s,(u,\sigma)) = \tau_x(\cdot\,|s,u,\sigma,q)T_q(q|s,u),\ q \in \mathcal{Q}, \tag{2.1}$$

which assigns to each $s \in \mathcal{S}$, $(u,\sigma) \in \mathcal{U} \times \Sigma$ a probability measure on the Borel space $(\mathcal{S}, \mathcal{B}(\mathcal{S}))$. Then, the DTSHS algorithm in Definition 16 can be rewritten in a more compact form as:

Algorithm 4 (DTSHS Execution - Markovian controls case)**.**
set $\mathbf{s}(0) = s_0$ *and* $k = 0$*;*

while $k < N$ *do*

set $(u_k, \sigma_k) = \mu_k(s_k)$*;*

extract from $\mathcal{S}$ *a value* s_{k+1} *for* $\mathbf{s}(k+1)$ *according to* $T_s(\cdot\,|s_k,(u_k,\sigma_k))$*;*

$k \to k+1$*;*

end.

□

This shows that a DTSHS $\mathscr{S} = (\mathcal{Q}, n, \mathcal{U}, \Sigma, T_x, T_q, R)$ can be described as a

controlled Markov process with state space $\mathcal{S} = \bigcup_{q\in\mathcal{Q}}\{q\} \times \mathbb{R}^{n(q)}$, control space $\mathcal{A} := \mathcal{U} \times \Sigma$, and controlled transition probability function $T_s : \mathcal{B}(\mathcal{S}) \times \mathcal{S} \times \mathcal{A} \to [0, 1]$ defined in (2.1), [Puterman, 1994]. This will be referred to in the following as *embedded controlled Markov process*.

As a consequence of this representation of $\mathscr{S}$, the execution $\{\mathbf{s}(k) = (\mathbf{q}(k), \mathbf{x}(k)), k \in [0, N]\}$ associated with $s_0 \in \mathcal{S}$ and $\mu \in \mathcal{M}_m$ is a stochastic process defined on the canonical sample space $\Omega = \mathcal{S}^{N+1}$,[1] endowed with its product topology $\mathcal{B}(\Omega)$, with probability measure $\mathcal{P}_\pi^\mu$ uniquely defined by the transition kernel T_s, the policy $\mu \in \mathcal{M}_m$, and the initial condition $s_0 \in \mathcal{S}$ [Bertsekas and Shreve, 1996, Proposition 7.45]. From the embedded Markov process representation of a DTSHS it also follows that the execution of a DTSHS associated with a Markov policy $\mu \in \mathcal{M}_m$ and an initial condition s_0 samples from the distribution π is an inhomogeneous Markov process with one-step transition kernels $T_s(\cdot|s, \mu_k(s))$, $k = 0, 1, \ldots, N-1$. In the following, we shall use the more compact notation $T_s^{\mu_k}(\cdot|s)$ for $T_s(\cdot|s, \mu_k(s))$.

[1] More precisely, the space of $N+1$-dimensional càdlàg paths with values in $\mathcal{S}$.

2.2 Theory

In this part, after framing the mathematical setup that will support the study, let us give a formal definition of the concept of probabilistic reachability, and its closely related concept of stochastic safety. By formulating an optimal control problem to maximize a safety probability criterion, we shall give precise procedures for solving this problem. We shall then extend the approach to the infinite-horizon scenario, and introduce related problems from Control Theory that are prone to be solved with the same technique. In particular, we discuss ways to embed the concept of performance in the above theoretical and computational framework.

2.2.1 Definition of the Concept

Let us consider the problem of determining the probability that the state of a DTSHS $\mathscr{S}$ will remain within a certain 'safe' set during a time horizon $[0, N]$, starting from an initial probability distribution π, under some control policy $\mu \in \mathcal{M}_m$. This probabilistic safety problem can be clearly classified as a stochastic reachability analysis problem.

Let the Borel set $A \in \mathcal{B}(S)$ represent a safe set. Our goal is setting up a reachability computation procedure to determine the probability that the execution associated with the Markov policy $\mu \in \mathcal{M}_m$ and initialization π will remain within A during the time horizon $[0, N]$:

$$\mathrm{p}_{\pi}^{\mu}(A) := \mathcal{P}_{\pi}^{\mu}\{\mathbf{s}(k) \in A \text{ for all } k \in [0, N]\}. \tag{2.2}$$

When π is concentrated at $\{s_0\}$, $s_0 \in \mathcal{S}$, we shall write simply $\mathcal{P}_{s_0}^{\mu}$.

If $\mathrm{p}_{\pi}^{\mu}(A) \geq 1 - \epsilon$, $\epsilon \in [0, 1]$, let us say that the system initialized at π is safe with at least probability $1 - \epsilon$ under policy μ.

Different initial conditions are characterized by different values of the probability $\mathrm{p}_{\pi}^{\mu}(A)$. Fix $\epsilon \in [0, 1]$. Let us define as *probabilistic safe set* with safety level $1 - \epsilon$ under policy μ the set

$$S^{\mu}(\epsilon) = \{s_0 \in \mathcal{S} : \mathrm{p}_{s}^{\mu}(A) \geq 1 - \epsilon\}, \tag{2.3}$$

of those initial conditions s_0 that correspond to a probability $\mathrm{p}_{\pi}^{\mu}(A)$ of remaining within the safe set A that is greater than or equal to $1-\epsilon$.

If for any initial condition $s_0 \in \mathcal{S}$ the control policy $\mu \in \mathcal{M}_m$ can be selected so as to maximize the probability of staying within A, then, let us define the set

$$S^{\star}(\epsilon) = \{s_0 \in \mathcal{S} : \sup_{\mu \in \mathcal{M}_m} \mathrm{p}_{s_0}^{\mu}(A) \geq 1-\epsilon\}. \tag{2.4}$$

By comparing the expressions for $S^{\mu}(\epsilon)$ and $S^{\star}(\epsilon)$, it is easily seen that $S^{\mu}(\epsilon) \subseteq S^{\star}(\epsilon)$, for each $\mu \in \mathcal{M}_m$ and for any $\epsilon \in [0,1]$, since in fact we are exploiting the best available control to achieve the ϵ-dependent reachability specification for the largest possible subset of the hybrid state space. The set $S^{\star}(\epsilon)$ is named the *maximal probabilistic safe set* with safety level $1-\epsilon$. Computing $S^{\star}(\epsilon)$ involves solving an optimization problem, and is a more challenging goal than computing $\mathrm{p}_{s_0}^{\mu}(A)$ and $S^{\mu}(\epsilon)$.

Note that the probability $\mathrm{p}_{\pi}^{\mu}(A)$ defined in (2.2) can be expressed as

$$\mathrm{p}_{\pi}^{\mu}(A) = 1 - \mathrm{P}_{\pi}^{\mu}(\bar{A}), \tag{2.5}$$

where $\bar{A}$ is the complement of A in $\mathcal{S}$. A similar expression intuitively holds if the initial distribution is condensed into a single point. Furthermore,

$$\mathrm{P}_{\pi}^{\mu}(\bar{A}) := \mathcal{P}_{\pi}^{\mu}\{\mathbf{s}(k) \in \bar{A} \text{ for some } k \in [0,N]\} \tag{2.6}$$

is the probability of entering $\bar{A}$ during the time interval $[0,N]$, given an initialization π. This leads to the following alternative expressions for $S^{\mu}(\epsilon)$ and $S^{\star}(\epsilon)$:

$$S^{\pi}(\epsilon) = \{s_0 \in \mathcal{S} : \mathrm{P}_{s_0}^{\mu}(\bar{A}) \leq \epsilon\} \tag{2.7}$$

$$S^{\star}(\epsilon) = \{s_0 \in \mathcal{S} : \inf_{\mu \in \mathcal{M}_m} \mathrm{P}_{s_0}^{\mu}(\bar{A}) \leq \epsilon\}. \tag{2.8}$$

In the following, let us show that

1. the problem of computing $\mathrm{p}_{\pi}^{\mu}(A)$, $\mathrm{P}_{\pi}^{\mu}(\bar{A})$, and $S^{\mu}(\epsilon)$ for $\mu \in \mathcal{M}_m$ can be solved by using a backward iterative procedure; and that

2. the problem of computing $S^{\star}(\epsilon)$ can be reduced to that of solving an optimal

control problem by dynamic programming.

These results are obtained by representing $\mathrm{p}^{\mu}_{\pi}(A)$ as a multiplicative cost function, and $\mathrm{P}^{\mu}_{\pi}(\bar{A})$ as a max cost function.

Let $\mathbf{1}_C : \mathcal{S} \to \{0,1\}$ denote the indicator function of a Borel set $C \in \mathcal{B}(\mathcal{S})$:

$$\mathbf{1}_C(s) = \begin{cases} 1, & \text{if } s \in C, \\ 0, & \text{if } s \notin C. \end{cases}$$

Multiplicative Cost. Observe that

$$\prod_{k=0}^{N} \mathbf{1}_A(s_k) = \begin{cases} 1, & \text{if } s_k \in A \text{ for all } k \in [0,N] \\ 0, & \text{otherwise,} \end{cases}$$

where $s_k \in \mathcal{S}$, $k \in [0,N]$ are a sequence of $N+1$ hybrid points—they may be thought of as being a single realization of a hybrid execution. Then, given a $\mu \in \mathcal{M}_m$ and an initial probability distribution π, $\mathrm{p}^{\mu}_{s_0}(A)$ in (2.2) can be expressed as the expectation with respect to the probability measure $\mathcal{P}^{\mu}_{s_0}$ of the Bernoulli random variable $\prod_{k=0}^{N} \mathbf{1}_A(\mathbf{s}(k))$:

$$\mathrm{p}^{\mu}_{s_0}(A) = \mathcal{P}^{\mu}_{s_0}\left(\prod_{k=0}^{N} \mathbf{1}_A(\mathbf{s}(k)) = 1\right) = E^{\mu}_{s_0}\left[\prod_{k=0}^{N} \mathbf{1}_A(\mathbf{s}(k))\right]. \tag{2.9}$$

We have implicitly assumed that $s_0 \in \mathcal{S}$ belongs to the support of π. From this expression it follows that

$$\mathrm{p}^{\mu}_{\pi}(A) = \int_{\mathcal{S}} E^{\mu}_{s_0}\big[\prod_{k=0}^{N} \mathbf{1}_A(\mathbf{s}(k))\big]\pi(ds), \tag{2.10}$$

where the conditional mean $E^{\mu}_{s_0}\big[\prod_{k=0}^{N} \mathbf{1}_A(\mathbf{s}(k))\big] = E^{\mu}_{\pi}[\prod_{k=0}^{N} \mathbf{1}_A(\mathbf{s}(k))|\, \mathbf{s}(0) = s]$ is well defined over the support of the probability measure π representing the distribution of $\mathbf{s}(0)$.

Max Cost. Proceeding similarly as in the above paragraph, observe that

$$\max_{k\in[0,N]} \mathbf{1}_{\bar{A}}(s_k) = \begin{cases} 1, & \text{if } s_k \in \bar{A} \text{ for some } k \in [0,N] \\ 0, & \text{otherwise,} \end{cases}$$

where $s_k \in \mathcal{S}$, $k \in [0,N]$ are a sequence of $N+1$ hybrid points—they may be thought of as being a single realization of a hybrid execution. Then, the probability $\mathrm{P}^{\mu}_{\pi}(\bar{A})$ in (2.6) can be expressed via

$$\mathrm{P}^{\mu}_{s_0}(\bar{A}) = \mathcal{P}^{\mu}_{s_0}\left(\max_{k\in[0,N]} \mathbf{1}_{\bar{A}}(\mathbf{s}(k)) = 1\right) = E^{\mu}_{s_0}\left[\max_{k\in[0,N]} \mathbf{1}_{\bar{A}}(\mathbf{s}(k))\right]. \tag{2.11}$$

From this expression it follows that

$$\mathrm{P}^{\mu}_{\pi}(\bar{A}) = \int_{\mathcal{S}} E^{\mu}_{s_0}\big[\max_{k\in[0,N]} \mathbf{1}_{\bar{A}}(\mathbf{s}(k))\big]\pi(ds), \tag{2.12}$$

where the conditional mean $E^{\mu}_{s_0}\big[\max_{k\in[0,N]} \mathbf{1}_{\bar{A}}(\mathbf{s}(k))\big] = E^{\mu}_{\pi}[\max_{k\in[0,N]} \mathbf{1}_{\bar{A}}(\mathbf{s}(k))|\, \mathbf{s}(0) = s]$ is well defined over the support of the probability measure π representing the distribution of $\mathbf{s}(0)$.

2.2.2 Probabilistic Reachability Computations

Let us next show how to compute $\mathrm{p}^{\mu}_{\pi}(A)$ and $\mathrm{P}^{\mu}_{\pi}(\bar{A})$ through a backward iterative procedure. To this purpose, let us recall that a Markov policy $\mu \in \mathcal{M}_m$ is a sequence $\mu = (\mu_0, \mu_1, \mu_2, \ldots, \mu_{N-1})$ of maps $\mu_l : \mathcal{S} \to \mathcal{U} \times \Sigma$, $l = 0, 1, 2, \ldots, N-1$.

Multiplicative Cost. Define the set of functions $V^{\mu}_k : \mathcal{S} \to [0,1]$, $k = 0, 1, \ldots, N$, as follows:

$$\begin{aligned} V^{\mu}_N(s) =& \mathbf{1}_A(s) \\ V^{\mu}_k(s) =& \mathbf{1}_A(s) \int_{\mathcal{S}^{N-k}} \prod_{l=k+1}^{N} \mathbf{1}_A(s_l) \prod_{l=k+1}^{N-1} T^{\mu_l}_s(ds_{l+1}|s_l)\, T^{\mu_k}_s(ds_{k+1}|s), \end{aligned} \tag{2.13}$$

$k = 0, 1, \ldots, N-1$, $s \in \mathcal{S}$. Recall that $T_s^{\mu_h}(\cdot|s_h)$ stands for $T_s(\cdot|s_h, \mu_h(s_h))$. The maps $T_s^{\mu_h}(\cdot|s_h)$, $h = 0, 1, \ldots, N-1$, are the one-step transition kernels of the embedded Markov process obtained by applying the Markov policy $\mu = (\mu_0, \mu_1, \ldots)$ to the DTSHS (see section 2.1.1). Then, it is easily seen that, by (2.9), if s belongs to the support of π, then $E_\pi^\mu \left[\prod_{l=k}^N \mathbf{1}_A(\mathbf{s}(l))|\mathbf{s}(k) = s\right]$ is well defined and equal to the right-hand side of (2.13), so that

$$V_k^\mu(s) = E_s^\mu \left[\prod_{l=k}^N \mathbf{1}_A(\mathbf{s}(l))\right], \; s \in \mathcal{S},$$

denotes the probability of remaining in A during the residual time horizon $[k, N]$, starting from s at time k. Furthermore, $V_0^\mu(s)$ evaluated at $s = s_0$ returns $\mathrm{p}_{s_0}^\mu(A)$. Hence, by (2.10) and the above expression, it is possible to explicitly express $\mathrm{p}_\pi^\mu(A)$ as

$$\mathrm{p}_\pi^\mu(A) = \int_{\mathcal{S}} V_0^\mu(s)\pi(ds).$$

Moreover, the probabilistic safe set with safety level $1-\epsilon$, $\epsilon \in [0, 1]$, according to (2.3), can be expressed in terms of function V_0^μ as follows:

$$S^\mu(\epsilon) = \{s_0 \in \mathcal{S} : V_0^\mu(s_0) \geq 1 - \epsilon\}.$$

Let $\mathcal{F}$ denote the set of functions from $\mathcal{S}$ to $\mathbb{R}$, and define the operator $H : \mathcal{S} \times \mathcal{U} \times \Sigma \times \mathcal{F} \to \mathbb{R}$ as follows:

$$\begin{aligned} H(s, (u, \sigma), Z) \;&:=\; T_q(q|s, u) \int_{\mathbb{R}^{n(q)}} Z((q, v))T_x(dv|s, u) \\ &+ \sum_{q' \neq q} T_q(q'|s, u) \int_{\mathbb{R}^{n(q')}} Z((q', v))R(dv|s, \sigma, q'), \end{aligned} \tag{2.14}$$

for any $s = (q, x) \in \mathcal{S}$, $(u, \sigma) \in \mathcal{U} \times \Sigma$, and $Z \in \mathcal{F}$. The operator H is easily seen to be a linear operator. Moreover, H applied to a constant function $\bar{Z}(s) = c$, $s \in \mathcal{S}$, returns the constant c for any value of the other arguments s and (u, σ): $H(s, (u, \sigma), \bar{Z}) = c$, $\forall s \in \mathcal{S}$, $(u, \sigma) \in \mathcal{U} \times \Sigma$. This is because $H(s, (u, \sigma), Z)$ is the integral over $\mathcal{S}$ of function Z with respect to the (conditional) probability measure

$T_s(\cdot|s,(u,\sigma))$ defined in (2.1).

With an argument inspired by a similar line of reasoning in [Kumar and Varaiya, 1986] for additive costs, let us prove the following lemma.

Lemma 1. *Fix a Markov policy $\mu = (\mu_0, \mu_1, \ldots, \mu_{N-1}) \in \mathcal{M}_m$. The functions $V_k^\mu : \mathcal{S} \to [0,1]$, $k = 0, 1 \ldots, N-1$, can be computed by the backward recursion:*

$$V_k^\mu(s) = \boldsymbol{1}_A(s) H(s, \mu_k(s), V_{k+1}^\mu),\ s \in \mathcal{S}, \tag{2.15}$$

initialized with $V_N^\mu(s) = \boldsymbol{1}_A(s)$, $s \in \mathcal{S}$. □

Proof. Let us start by observing that, given the definition of T_s in (2.1) in terms of its components and that of H in (2.14), equation (2.15) can be rewritten as

$$V_k^\mu(s) = \mathbf{1}_A(s) \int_{\mathcal{S}} V_{k+1}^\mu(s_{k+1}) T_s(ds_{k+1}|s, \mu_k(s)).$$

From the expression in (2.13) of V_k^μ, it holds true that

$$\begin{aligned} V_{N-1}^\mu(s) &= \mathbf{1}_A(s) \int_{\mathcal{S}} \mathbf{1}_A(s_N) T_s^{\mu_{N-1}}(ds_N|s) \\ &= \mathbf{1}_A(s) \int_{\mathcal{S}} V_N^\mu(s_N) T_s(ds_N|s, \mu_{N-1}(s)) \end{aligned}$$

so that equation (2.15) is proven for $k = N-1$. For $k < N-1$, V_k^μ can be expanded as follows

$$\begin{aligned} V_k^\mu(s) &= \mathbf{1}_A(s) \int_{\mathcal{S}} \mathbf{1}_A(s_{k+1}) \Big(\int_{\mathcal{S}^{N-k-1}} \prod_{l=k+2}^{N} \mathbf{1}_A(s_l) \\ &\qquad \prod_{l=k+2}^{N-1} T_s^{\mu_l}(ds_{l+1}|s_l) T_s^{\mu_{k+1}}(ds_{k+2}|s_{k+1}) \Big) T_s^{\mu_k}(ds_{k+1}|s) \\ &= \mathbf{1}_A(s) \int_{\mathcal{S}} V_{k+1}^\mu(s_{k+1}) T_s^{\mu_k}(ds_{k+1}|s), \end{aligned}$$

which concludes the proof. □

Max Cost. Define the set of functions $W_k^\pi : \mathcal{S} \to [0,1]$, $k = 0, 1, \ldots, N$, as follows:

$$\begin{aligned} W_N^\mu(s) =& \mathbf{1}_{\bar{A}}(s) \\ W_k^\mu(s) =& \mathbf{1}_{\bar{A}}(s) + \mathbf{1}_A(s) \int_{\mathcal{S}^{N-k}} \max_{l \in [k+1,N]} \mathbf{1}_{\bar{A}}(s_l) \prod_{l=k+1}^{N-1} T_s^{\mu_l}(ds_{l+1}|s_l) T_s^{\mu_k}(ds_{k+1}|s), \end{aligned} \quad (2.16)$$

$k = 0, 1, \ldots, N-1$, $s \in \mathcal{S}$. Given a properly sized policy μ, the quantity $W_k^\mu(s)$ expresses the probability of exiting A for the residual time horizon $[k, N]$, starting from $s \in \mathcal{S}$. Given the expression of $\mathrm{P}_\pi^\mu(\bar{A})$ as a max cost in (2.11), it is easy to show that $W_0^\mu(s)$ evaluated at $s = s_0$ returns $\mathrm{P}_{s_0}^\mu(\bar{A})$ since

$$W_0^\mu(s_0) = E_{s_0}^\mu \left[\max_{l \in [0,N]} \mathbf{1}_{\bar{A}}(\mathbf{s}(l)) \right].$$

Again, the above expression assumes that s_0 belongs to the support of π. Furthermore, by (2.12) and the above expression, it is possible to express $\mathrm{P}_\pi^\mu(\bar{A})$ as

$$\mathrm{P}_\pi^\mu(\bar{A}) = \int_{\mathcal{S}} W_0^\mu(s) \pi(ds).$$

Also, based on (2.7), the probabilistic safe set with safety level $1 - \epsilon$, $\epsilon \in [0,1]$, can be expressed in terms of W_0^μ as

$$S^\mu(\epsilon) = \{s_0 \in \mathcal{S} : W_0^\mu(s_0) \leq \epsilon\}.$$

From the definition of W_k^μ in (2.16), and that of H in (2.14), Lemma 2 follows:

Lemma 2. *Fix a Markov policy $\mu = (\mu_0, \mu_1, \ldots, \mu_{N-1}) \in \mathcal{M}_m$. The functions $W_k^\mu : \mathcal{S} \to [0,1]$, $k = 0, 1 \ldots, N-1$, can be computed by the backward recursion:*

$$W_k^\mu(s) = \boldsymbol{1}_{\bar{A}}(s) + \boldsymbol{1}_A(s) H(s, \mu_k(s), W_{k+1}^\mu), \; s \in \mathcal{S}, \quad (2.17)$$

initialized with $W_N^\mu(s) = \boldsymbol{1}_{\bar{A}}(s)$, $s \in \mathcal{S}$. □

Proof. From the Definition(2.16) of W_k^μ, it holds that

$$\begin{aligned} W_{N-1}^\mu(s) &= \mathbf{1}_{\bar{A}}(s) + \mathbf{1}_A(s) \int_{\mathcal{S}} \mathbf{1}_{\bar{A}}(s_N) T_s^{\mu_{N-1}}(ds_N|s) \\ &= \mathbf{1}_{\bar{A}}(s) + \mathbf{1}_A(s) \int_{\mathcal{S}} W_N^\mu(s_N) T_s^{\mu_{N-1}}(ds_N|s) \end{aligned}$$

so that equation (2.17) is proven for $k = N-1$. For $k < N-1$, W_k^μ can be expanded as follows

$$\begin{aligned} W_k^\mu(s) = \mathbf{1}_{\bar{A}}(s) + \mathbf{1}_A(s) \int_{\mathcal{S}} \Big(& \mathbf{1}_{\bar{A}}(s_{k+1}) + \mathbf{1}_A(s_{k+1}) \\ & \int_{\mathcal{S}^{N-k-1}} \max_{l \in [k+2,N]} \mathbf{1}_{\bar{A}}(s_l) \prod_{l=k+2}^{N-1} T_s^{\mu_l}(ds_{l+1}|s_l) T_s^{\mu_{k+1}}(ds_{k+2}|s_{k+1}) \Big) T_s^{\mu_k}(ds_{k+1}|s) \\ = \mathbf{1}_{\bar{A}}(s) + \mathbf{1}_A(s) & \int_{\mathcal{S}} W_{k+1}^\mu(s_{k+1}) T_s^{\mu_k}(ds_{k+1}|s) \end{aligned}$$

which concludes the proof. □

It is worth noting that the iterative backward recursion derived in Lemma 2 is much like that in [Prandini and Hu, 2006a; Prandini and Hu, 2006b] for reachability computations on the Markov chain approximation of certain classes of *uncontrolled* continuous time stochastic hybrid systems.

Equivalence of the two Representations. Since, for any $s_l \in \mathcal{S}, l = 0, 1, \ldots, N$, $\prod_{l=k}^{N} \mathbf{1}_A(s_l) = 1 - \max_{l \in [k,N]} \mathbf{1}_{\bar{A}}(s_l)$, $k = 0, 1, \ldots, N$, not surprisingly, the following equivalence result holds.

Lemma 3. *Fix a Markov policy $\mu = (\mu_0, \mu_1, \ldots, \mu_{N-1}) \in \mathcal{M}_m$. Then, for any $k \in [0, N]$, $W_k^\mu(s) = 1 - V_k^\mu(s)$, $s \in \mathcal{S}$.* □

Proof. The statement trivially holds for $k = N$. Assume now that it holds at time

$k+1$. Then,

$$\begin{aligned}
W_k^\mu(s) =& \text{[by Lemma 2]} \\
=& \mathbf{1}_{\bar{A}}(s) + \mathbf{1}_A(s) H(s, \mu_k(s), W_{k+1}^\mu) \\
=& \text{[by induction assumption]} \\
=& 1 - \mathbf{1}_A(s) + \mathbf{1}_A(s) H(s, \mu_k(s), 1 - V_{k+1}^\mu) \\
=& \text{[by the properties of the operator } H\text{]} \\
=& 1 - \mathbf{1}_A(s) + \mathbf{1}_A(s)(1 - H(s, \mu_k(s), V_{k+1}^\mu)) \\
=& 1 - \mathbf{1}_A(s) H(s, \mu_k(s), V_{k+1}^\mu)) \\
=& \text{[by Lemma 1]} \\
=& 1 - V_k^\mu(s),\ s \in \mathcal{S},
\end{aligned}$$

so that the statement holds for any $k = 0, 1, \ldots, N$. □

2.2.3 Maximal Probabilistic Safe Sets Computation

The calculation of the maximal probabilistic safe set $S^\star(\epsilon)$ defined in (2.4) (or equivalently in (2.8)) amounts to finding the supremum over the Markov policies of the probability $\mathrm{p}_\pi^\mu(A)$ of remaining within the safe set A starting from s_0, for all s_0 inside A (the probability of remaining within A starting from $s_0 \notin A$ is 0 for any policy). A policy that achieves this supremum is said to be *maximally safe*. More precisely,

Definition 17 (Maximally safe policy)**.** *Let $\mathscr{S}$ be a DTSHS, and $A \in \mathcal{B}(\mathcal{S})$ a safe set. A Markov policy $\mu^\star$ is maximally safe if $\mathrm{p}_{s_0}^{\mu^\star}(A) = \sup_{\mu \in \mathcal{M}_m} \mathrm{p}_{s_0}^\mu(A)$, $\forall s_0 \in A$.* □

Note that, in view of Lemma 3, a maximally safe policy can be equivalently characterized as that policy $\mu^\star \in \mathcal{M}_m$ that achieves the minimum over A of $\mathrm{P}_{s_0}^\mu(\bar{A})$: $\mathrm{P}_{s_0}^{\mu^\star}(\bar{A}) = \inf_{\mu \in \mathcal{M}_m} \mathrm{P}_{s_0}^\mu(\bar{A})$, $\forall s_0 \in A$.

In general, a maximally safe policy is not guaranteed to exist. Let us next provide sufficient conditions for the existence of a maximally safe Markov policy, and describe an algorithm to compute $\sup_{\mu \in \mathcal{M}_m} \mathrm{p}_{s_0}^\mu(A)$ in terms of the multiplicative cost, and $\inf_{\mu \in \mathcal{M}_m} \mathrm{P}_{s_0}^\mu(\bar{A})$ in terms of the max cost.

Multiplicative Cost. Let us now show how to compute a maximally safe Markov policy $\mu^\star \in \mathcal{M}_m$ through a recursion much like that in Lemma 1, based on the representation (2.9) of $\mathrm{p}^\mu_{s_0}(A)$ as a multiplicative cost. The proof is inspired by [Bertsekas and Shreve, 1996, section 11.3], addressing a finite horizon stochastic optimal control problem with a multiplicative cost to be minimized.

Theorem 5. *Define $V^\star_k : \mathcal{S} \to [0,1]$, $k = 0, 1, \ldots, N$, by the recursion:*

$$V^\star_k(s) = \sup_{(u,\sigma)\in\mathcal{U}\times\Sigma} \boldsymbol{1}_A(s) H(s,(u,\sigma),V^\star_{k+1}), \tag{2.18}$$

$s \in \mathcal{S}$, initialized with $V^\star_N(s) = \boldsymbol{1}_A(s)$, $s \in \mathcal{S}$.
Then, $V^\star_0(s_0) = \sup_{\mu\in\mathcal{M}_m} \mathrm{p}^\mu_{s_0}(A)$, $s_0 \in \mathcal{S}$.
If $\mu^\star_k : \mathcal{S} \to \mathcal{U}\times\Sigma, k \in [0, N-1]$, is such that

$$\mu^\star_k(s) = \arg \sup_{(u,\sigma)\in\mathcal{U}\times\Sigma} H(s,(u,\sigma),V^\star_{k+1}),\ s \in A, \tag{2.19}$$

then, $\mu^\star = (\mu^\star_0, \mu^\star_1, \ldots, \mu^\star_{N-1})$ is a maximally safe Markov policy. A sufficient condition for the existence of such a $\mu^\star$ is that $U_k(s,\lambda) = \{(u,\sigma) \in \mathcal{U}\times\Sigma :\ H(s,(u,\sigma),V^\star_{k+1}) \geq \lambda\}$ is compact for all $s \in A$, $\lambda \in \mathbb{R}$, $k \in [0, N-1]$. □

Proof. For ease of reference to [Bertsekas and Shreve, 1996, section 11.3], let us set $J^\mu_k := -V^\mu_{N-k}$, $\mu \in \mathcal{M}_m$, and $J^\star_k := -V^\star_{N-k}$, $k = 0, 1, \ldots, N$, and rewrite equation (2.19) and the recursions (2.15) and (2.18) in terms of these functions as:

$$\mu^\star_k(s) = \arg \inf_{(u,\sigma)\in\mathcal{U}\times\Sigma} H(s,(u,\sigma),J^\star_{N-k-1}),\ s \in A, \tag{2.20}$$

$$J^\mu_k(s) = \mathbf{1}_A(s) H(s,\mu_{N-k}(s),J^\mu_{k-1}) \tag{2.21}$$

$$J^\star_k(s) = \inf_{(u,\sigma)\in\mathcal{U}\times\Sigma} \mathbf{1}_A(s) H(s,(u,\sigma),J^\star_{k-1}), \tag{2.22}$$

initialized with $J^\mu_0(s) = J^\star_0(s) = -\mathbf{1}_A(s)$, $s \in \mathcal{S}$.

Consider a (universally measurable) function $\mu : \mathcal{S} \to \mathcal{U}\times\Sigma$ and define the map $T_\mu : \mathcal{F} \to \mathcal{F}$ as $T_\mu[J](s) = K(s,\mu(s),J)$, $s \in \mathcal{S}$, where $K(s,(u,\sigma),J) = \mathbf{1}_A(s) H(s,(u,\sigma),J)$, $s \in \mathcal{S}$, $(u,\sigma) \in \mathcal{U}\times\Sigma$, $J \in \mathcal{F}$.

Let $\tilde{\mathcal{F}} \subset \mathcal{F}$ denote the set of universally measurable real functions $J : \mathcal{S} \to \mathbb{R}$. The map T_μ preserves the universal measurability property: if $J \in \tilde{\mathcal{F}}$, then, $T_\mu[J] \in \tilde{\mathcal{F}}$. This is because the integration of a universally measurable function with respect to the stochastic kernel involved in the computation of $H(s, \mu(s), J)$ (see (2.14)) is a universally measurable function, and its product with the Borel measurable function $\mathbf{1}_A(s)$ remains universally measurable.

Observe that, since the recursion (2.21) can be rewritten as $J_k^\mu = T_{\mu_{N-k}}[J_{k-1}^\mu]$ and $J_0^\mu \in \tilde{\mathcal{F}}$, we then have that $J_k^\mu \in \tilde{\mathcal{F}}$, $k = 1, 2, \ldots, N$.

Map T_μ satisfies also the following properties: for all $J, J' \in \tilde{\mathcal{F}}$ such that $J(s) \leq J'(s)$, $\forall s \in \mathcal{S}$, then $T_\mu[J](s) \leq T_\mu[J'](s)$, $\forall s \in \mathcal{S}$ (monotonicity, [Bertsekas and Shreve, 1996, section 6.2]), and for any $J \in \tilde{\mathcal{F}}$ and any real number $r > 0$,

$$T_\mu[J](s) \leq T_\mu[J + r](s) \leq T_\mu[J](s) + r, \; s \in \mathcal{S}. \tag{2.23}$$

The monotonicity property immediately follows from the definition of T_μ. As for property (2.23), it is easily shown observing that, by the definition of K and the properties of H, the following chain of equalities hold: $K(s, (u, \sigma), J + r) = \mathbf{1}_A(s) H(s, (u, \sigma), J + r) = \mathbf{1}_A(s) H(s, (u, \sigma), J) + \mathbf{1}_A(s) r$, $s \in \mathcal{S}$, $(u, \sigma) \in \mathcal{U} \times \Sigma$, and, hence, given that $\mathbf{1}_A(s)$ is either equal to 0 or to 1, $K(s, (u, \sigma), J) \leq K(s, (u, \sigma), J + r) \leq K(s, (u, \sigma), J) + r$, $s \in \mathcal{S}$, $(u, \sigma) \in \mathcal{U} \times \Sigma$.

Now, define the map $T : \mathcal{F} \to \mathcal{F}$ as $T[J](s) = \inf_{(u,\sigma) \in \mathcal{U} \times \Sigma} K(s, (u, \sigma), J)$, $s \in \mathcal{S}$. Then, the recursion (2.22) can be rewritten as $J_k^\star = T[J_{k-1}^\star]$, and, from this latter expression, it follows that $J_k^\star = T^k[J_0^\star]$, $k = 0, 1, \ldots, N$, where $T^0[J] = J$ and $T^k[J] = T[T^{k-1}[J]]$. Let $\mathcal{F}^* \subset \tilde{\mathcal{F}}$ denote the set of lower-semianalytic functions. The map T preserves the lower-semianalytic property: if $J \in \mathcal{F}^*$, then, $T[J] \in \mathcal{F}^*$. This follows from [Bertsekas and Shreve, 1996, Proposition 7.47], given that $H(s, (u, \sigma), J)$ is lower-semianalytic as a function of its arguments $s \in \mathcal{S}$ and $(u, \sigma) \in \mathcal{U} \times \Sigma$ [Bertsekas and Shreve, 1996, Proposition 7.48], and, hence, $K(s, (u, \sigma), J) = \mathbf{1}_A(s) H(s, (u, \sigma), J)$ is lower-semianalytic as well [Bertsekas and Shreve, 1996, Lemma 7.30(4)]. Since $J_k^\star = T[J_{k-1}^\star]$ and $J_0^\star \in \mathcal{F}^*$, we then have that $J_k^\star \in \mathcal{F}^*$, $k = 1, 2, \ldots, N$.

After these preliminary considerations, let us prove by induction on N the follow-

ing two statements:

$$1.\ \inf_{\mu} E_s^{\mu}[-\prod_{k=0}^{N}\mathbf{1}_A(s_k)] = T^N[J_0^\star](s),\ s\in\mathcal{S}$$

$$2.\ \forall\epsilon>0,\ \exists\,\mu_\epsilon=(\mu_{\epsilon,0},\mu_{\epsilon,1},\dots)\in\mathcal{M}_m:\forall s\in\mathcal{S},$$

$$\inf_{\mu} E_s^{\mu}[-\prod_{k=0}^{N}\mathbf{1}_A(s_k)] \le J_N^{\mu_\epsilon}(s) \le \inf_{\mu} E_s^{\mu}[-\prod_{k=0}^{N}\mathbf{1}_A(s_k)] + \epsilon.$$

Note that by the first statement, it follows that $V_0^\star(\pi) = -J_N^\star(\pi) = -T^N[J_0^\star](\pi) = \sup_{\mu\in\mathcal{M}_m} \mathrm{p}_\pi^\mu(A)$, for any $\pi\in\mathcal{S}$, so that the first part of the theorem is proven. As for the second statement, observe that for any $\mu_\epsilon\in\mathcal{M}_m$

$$J_N^{\mu_\epsilon}(s) = -V_0^{\mu_\epsilon}(s) = E_s^{\mu_\epsilon}[-\prod_{k=0}^{N}\mathbf{1}_A(s_k)] \ge \inf_{\mu} E_s^{\mu}[-\prod_{k=0}^{N}\mathbf{1}_A(s_k)],$$

so that the part of the second statement that needs to be actually proven is the right-hand-side.

The statements clearly hold for $N=0$. Suppose that they hold for $N=h$. This implies that $\forall\epsilon>0,\ \exists\,\mu_\epsilon=(\mu_{\epsilon,0},\mu_{\epsilon,1},\dots)\in\mathcal{M}_m$ such that

$$J_h^{\mu_\epsilon}(s) \le \inf_{\mu} E_s^{\mu}[-\prod_{l=0}^{h}\mathbf{1}_A(s_l)] + \epsilon, \quad s\in\mathcal{S}.$$

For any universally measurable function $\mu:\mathcal{S}\to\mathcal{U}\times\Sigma$, we then have that, by the monotonicity of T_μ and (2.23),

$$\begin{aligned} T_\mu[J_h^{\mu_\epsilon}](s) &\le T_\mu[\inf_{\mu} E_s^{\mu}[-\prod_{l=0}^{h}\mathbf{1}_A(s_l)] + \epsilon] \\ &\le T_\mu[\inf_{\mu} E_s^{\mu}[-\prod_{l=0}^{h}\mathbf{1}_A(s_l)]] + \epsilon, \quad s\in\mathcal{S}. \end{aligned}$$

Now, if a policy $\bar{\mu}_\epsilon = (\mu, \mu_{\epsilon,0}, \mu_{\epsilon,1}, \dots)$ is considered, it is easily seen that

$$\inf_\mu E_s^\mu[-\prod_{l=0}^{h+1} \mathbf{1}_A(s_l)] \leq J_{h+1}^{\bar{\mu}_\epsilon}(s) = T_\mu[J_h^{\mu_\epsilon}](s), \ s \in \mathcal{S},$$

which, combined with the inequality above, leads to:

$$\inf_\mu E_s^\mu[-\prod_{l=0}^{h+1} \mathbf{1}_A(s_l)] \leq T_\mu[\inf_\mu E_s^\mu[-\prod_{l=0}^{h} \mathbf{1}_A(s_l)]] + \epsilon, \ s \in \mathcal{S},$$

for any universally measurable function $\mu : \mathcal{S} \to \mathcal{U} \times \Sigma$. From this, it follows that

$$\begin{aligned}\inf_\mu E_s^\mu[-\prod_{l=0}^{h+1} \mathbf{1}_A(s_l)] &\leq T[\inf_\mu E_s^\mu[-\prod_{l=0}^{h} \mathbf{1}_A(s_l)]] \\ &= T^{h+1}[J_0^\star](s), \quad s \in \mathcal{S},\end{aligned}$$

where the last equality is due to the induction hypothesis. On the other hand, it is true that

$$T^{h+1}[J_0^\star](s) \leq \inf_\mu E_s^\mu[-\prod_{l=0}^{h+1} \mathbf{1}_A(s_l)], \quad s \in \mathcal{S},$$

which allows one to conclude that

$$\inf_\mu E_s^\mu[-\prod_{l=0}^{h+1} \mathbf{1}_A(s_l)] = T^{h+1}[J_0^\star](s), \quad s \in \mathcal{S}. \tag{2.24}$$

Let us now proceed with the second statement.
By the induction hypothesis, for any $\bar{\epsilon} > 0$ there exists $\bar{\mu} = (\bar{\mu}_0, \bar{\mu}_1, \dots) \in \mathcal{M}_m$ such that

$$J_h^{\bar{\mu}}(s) \leq \inf_\mu E_s^\mu[-\prod_{l=0}^{h} \mathbf{1}_A(s_l)] + \frac{\bar{\epsilon}}{2}, \quad s \in \mathcal{S}.$$

Also, by [Bertsekas and Shreve, 1996, Proposition 7.50] there exists a universally

measurable function $\bar{\mu} : \mathcal{S} \to \mathcal{U} \times \Sigma$ such that

$$T_{\bar{\mu}}[E_s^{\mu}[-\prod_{l=0}^{h} \mathbf{1}_A(s_l)]](s) \leq T[E_s^{\mu}[-\prod_{l=0}^{h} \mathbf{1}_A(s_l)]](s) + \frac{\bar{\epsilon}}{2},$$

$s \in \mathcal{S}$. Then, if $\mu_{\bar{\epsilon}} = (\bar{\mu}, \bar{\mu}_0, \bar{\mu}_1, \dots)$ is considered, by the monotonicity of $T_{\bar{\mu}}$ and (2.23), we obtain

$$\begin{aligned} J_{h+1}^{\mu_{\bar{\epsilon}}}(s) = T_{\bar{\mu}}[J_h^{\bar{\mu}}](s) &\leq T_{\bar{\mu}}[\inf_{\mu} E_s^{\mu}[-\prod_{l=0}^{h} \mathbf{1}_A(s_l)]] + \frac{\bar{\epsilon}}{2} \\ &\leq T[\inf_{\mu} E_s^{\mu}[-\prod_{l=0}^{h} \mathbf{1}_A(s_l)]] + \bar{\epsilon}, \quad s \in \mathcal{S}. \end{aligned}$$

By the induction hypothesis and (2.24), we finally get

$$J_{h+1}^{\mu_{\bar{\epsilon}}}(s) \leq T^{h+1}[J_0^{\star}] + \bar{\epsilon} = \inf_{\mu} E_s^{\mu}[-\prod_{l=0}^{h+1} \mathbf{1}_A(s_l)] + \bar{\epsilon},$$

$s \in \mathcal{S}$, which concludes the proof of the two statements.

Next, let us show that $\mu^{\star} = (\mu_0^{\star}, \mu_1^{\star}, \dots)$ satisfying (2.20) is a Markov policy and that it is maximally safe. To this purpose, note first that function $\mu_k^{\star}$ satisfying (2.20) can be characterized through the equation

$$T_{\mu_k^{\star}}[J_{N-k-1}^{\star}](s) = \inf_{(u,\sigma) \in \mathcal{U} \times \Sigma} K(s, (u, \sigma), J_{N-k-1}^{\star}) = J_{N-k}^{\star}(s), \quad s \in \mathcal{S}.$$

As discussed at the beginning of this proof, $J_{N-k}^{\star} \in \mathcal{F}^*$ and, hence, $K(s, (u, \sigma), J_{N-k}^{\star})$ is lower-semianalytic as a function of $s \in \mathcal{S}$ and $(u, \sigma) \in \mathcal{U} \times \Sigma$. Then, if its infimum with respect to $(u, \sigma) \in \mathcal{U} \times \Sigma$ is attained for any $s \in \mathcal{S}$ (for $s \in \bar{A}$ this is always the case), the resulting function $\mu_k^{\star} : \mathcal{S} \to \mathcal{U} \times \Sigma$ is universally measurable, [Bertsekas and Shreve, 1996, Proposition 7.50].

Now observe that

$$\begin{aligned}\inf_{\mu} E_s^{\mu}[-\prod_{l=0}^{N}\mathbf{1}_A(s_l)] = J_N^\star(s) &= T_{\mu_0^\star}[J_{N-1}^\star](s)\\ &= T_{\mu_0^\star}\left[T_{\mu_1^\star}[J_{N-2}^\star]\right](s) = \dots\\ &= T_{\mu_0^\star}[T_{\mu_1^\star}[\dots T_{\mu_{N-1}^\star}[J_0^\star]]](s)\\ &= J_N^{\mu^\star}(s) = E_s^{\mu^\star}[-\prod_{l=0}^{N}\mathbf{1}_A(s_l)],\end{aligned}$$

$s \in \mathcal{S}$, which shows that $\mu^\star$ is maximally safe.
For any $s \in \mathcal{S}$ and $k \in [0, N-1]$, a sufficient condition for the existence of a minimum over $\mathcal{U}\times\Sigma$ of function $K(s,(u,\sigma),J_{N-k-1}^\star)$ is that $Z_k(s,\alpha) = \{(u,\sigma) \in \mathcal{U}\times\Sigma :\ K(s,(u,\sigma),J_{N-k}-1^\star) \le \alpha\}$ is compact, [Bertsekas and Shreve, 1996, Lemma 3.1]. Since $J_{N-k-1}^\star = -V_{k+1}^\star$, then $K(s,(u,\sigma),J_{N-k-1}^\star) = \mathbf{1}_A(s)H(s,(u,\sigma),J_{N-k-1}^\star) = -\mathbf{1}_A(s)H(s,(u,\sigma),V_{k+1}^\star)$, from which the condition on $U_k(s,\lambda)$ easily follows. □

Max Cost. In the following theorem, let us describe an algorithm to compute a maximally safe Markov policy $\mu^\star \in \mathcal{M}_m$ based on the representation (2.11) of $\mathrm{P}_\pi^\mu(\bar{A})$ as a max cost, by a recursion much like that in Lemma 2.

Theorem 6. *Define $W_k^\star : \mathcal{S} \to [0,1]$, $k = 0,1,\dots,N$, by the recursion:*

$$W_k^\star(s) = \inf_{(u,\sigma)\in\mathcal{U}\times\Sigma}\left(\mathbf{1}_{\bar{A}}(s) + \mathbf{1}_A(s)H(s,(u,\sigma),W_{k+1}^\star)\right),$$

$s \in \mathcal{S}$, initialized with $W_N^\star(s) = \mathbf{1}_{\bar{A}}(s)$, $s \in \mathcal{S}$.
Then, $W_0^\star(\pi) = \inf_{\mu\in\mathcal{M}_m}\mathrm{P}_\pi^\mu(\bar{A})$, $\pi \in \mathcal{S}$.
If $\mu_k^\star : \mathcal{S} \to \mathcal{U}\times\Sigma$, $k \in [0, N-1]$, is such that

$$\mu_k^\star(s) = \arg\inf_{(u,\sigma)\in\mathcal{U}\times\Sigma} H(s,(u,\sigma),W_{k+1}^\star), \forall s \in A,$$

then, $\mu^\star = (\mu_0^\star, \mu_1^\star, \dots, \mu_{N-1}^\star, \dots)$ is a maximally safe Markov policy. A sufficient condition for the existence of such a $\mu^\star$ is that $U_k(s,\lambda) = \{(u,\sigma) \in \mathcal{U}\times\Sigma :\ H(s,(u,\sigma),W_{k+1}^\star) \le \lambda\}$ is compact for all $s \in A$, $\lambda \in \mathbb{R}$, $k \in [0, N-1]$. □

Proof. Let us start proving that for any $k \in [0, N]$,

$$W_k^\star(s) = 1 - V_k^\star(s),\ s \in \mathcal{S}. \tag{2.25}$$

The statement is trivially satisfied for $k = N$, since $W_N^\star(s) = \mathbf{1}_{\bar{A}}(s) = 1 - \mathbf{1}_A(s) = 1 - V_N^\star(s)$, $s \in \mathcal{S}$. Assume that it is valid for $k+1$. Then,

$$\begin{aligned}
W_k^\star(s) &= \inf_{(u,\sigma)\in\mathcal{U}\times\Sigma} (\mathbf{1}_{\bar{A}}(s) + \mathbf{1}_A(s) H(s,(u,\sigma),W_{k+1}^\star)) \\
&= [\text{by induction assumption}] \\
&= \inf_{(u,\sigma)\in\mathcal{U}\times\Sigma} (\mathbf{1}_{\bar{A}}(s) + \mathbf{1}_A(s) H(s,(u,\sigma),1 - V_{k+1}^\star)) \\
&= [\text{by the properties of the operator } H] \\
&= \inf_{(u,\sigma)\in\mathcal{U}\times\Sigma} (\mathbf{1}_{\bar{A}}(s) + \mathbf{1}_A(s)(1 - H(s,(u,\sigma),V_{k+1}^\star))) \\
&= 1 - \sup_{(u,\sigma)\in\mathcal{U}\times\Sigma} \mathbf{1}_A(s) H(s,(u,\sigma),V_{k+1}^\star) \\
&= [\text{by Theorem 5}] \\
&= 1 - V_k^\star(s), \quad s \in \mathcal{S}.
\end{aligned}$$

It then easily follows from Theorem 5 and the definitions of $\mathrm{p}_\pi^\mu(A)$ and $\mathrm{P}_\pi^\mu(\bar{A})$ that

$$\begin{aligned}
W_0^\star(\pi) = 1 - V_0^\star(\pi) &= 1 - \sup_{\mu\in\mathcal{M}_m} \mathrm{p}_\pi^\mu(A) \\
&= \inf_{\mu\in\mathcal{M}_m} (1 - \mathrm{p}_\pi^\mu(A)) = \inf_{\mu\in\mathcal{M}_m} \mathrm{P}_\pi^\mu(\bar{A}).
\end{aligned}$$

Furthermore, in view of the duality equation (2.25) the characterization through the $V_k^\star$ functions of a maximally safe policy $\mu_k^\star : \mathcal{S} \to \mathcal{U} \times \Sigma$, $k \in [0, N-1]$, in Theorem 5 can be equivalently expressed in terms of the $W_k^\star$ functions as follows:

$$\begin{aligned}
\mu_k^\star(s) &= \arg \sup_{(u,\sigma)\in\mathcal{U}\times\Sigma} H(s,(u,\sigma),V_{k+1}^\star) \\
&= \arg \sup_{(u,\sigma)\in\mathcal{U}\times\Sigma} H(s,(u,\sigma),1 - W_{k+1}^\star) \\
&= \arg \inf_{(u,\sigma)\in\mathcal{U}\times\Sigma} H(s,(u,\sigma),W_{k+1}^\star), \quad s \in A.
\end{aligned}$$

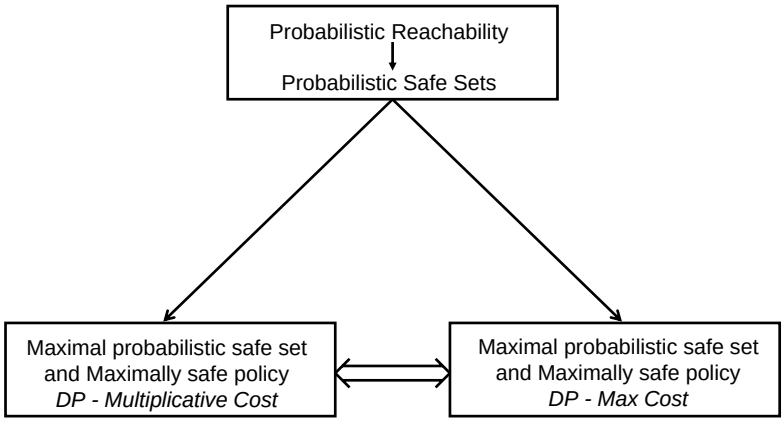

Figure 2.3: Dual interpretation of the probabilistic reachability and safety problem.

A sufficient condition on the control space to ensure the existence of this optimal argument at each time step is again easily derived from the corresponding one in Theorem 5. □

Remark 11. *If the control input spaces $\mathcal{U}$ and Σ are both finite sets, then a maximally safe policy is guaranteed to exist.* □

2.2.4 Extensions to the Infinite Horizon Case

Motivated by the quest for generality, as well as by the obtained outputs from the simulations of the case study (see in particular the outputs of Figure 2.9), let us extend the reachability analysis problem for DTSHS to the infinite time-horizon setting and address the question of convergence of the optimal control law to a stationary policy. The extension will adhere to the max cost approach.

Let us consider a system that is described by a controlled DTSHS model $\mathscr{S}$ (see Definition 13). The sets $\mathcal{M}$ and $\mathcal{M}_m$ of feedback and Markov policies are extension to the infinite horizon case of those introduced in section 2.1.1. A Markov policy $\mu \in \mathcal{M}_m$ is said to be stationary if $\mu = (\bar{\mu}, \bar{\mu}, \bar{\mu}, \dots)$, with $\bar{\mu} : \mathcal{S} \to \mathcal{U} \times \Sigma$ universally measurable. The execution of the DTSHS $\mathscr{S}$ associated with some policy μ and initial distribution π is easily obtained by extending Definition 16 to the infinite horizon.

Then, the execution $\{\mathbf{s}(k), k \geq 0\}$ associated with $\mu \in \mathcal{M}_m$ and π is a stochastic process defined on the canonical sample space $\Omega = \mathcal{S}^\infty$, endowed with its product topology $\mathcal{B}(\Omega)$, with probability measure $\mathcal{P}^\mu_{\pi,\infty}$ uniquely defined by the transition kernel T_s, the policy μ, and the initial distribution π (see [Bertsekas and Shreve, 1996, Proposition 7.45]).

For a given policy $\mu \in \mathcal{M}_m$ and initial distribution π, let

$$\mathrm{P}^\mu_{\pi,\infty}(\bar{A}) := \mathcal{P}^\mu_{\pi,\infty}(\mathbf{s}(k) \in \bar{A} \text{ for some } k \geq 0),$$

be the probability of entering the unsafe region specified by $\bar{A} \in \mathcal{B}(\mathcal{S})$. If π is concentrated in a single point s, let us use the notation $\mathrm{P}^\mu_{s,\infty}(\bar{A})$.

The goal is again that of finding an optimal Markov policy that singles out the *maximal probabilistic safe set*

$$S^\star_\infty(\epsilon) := \{s_0 \in \mathcal{S} : \inf_{\mu \in \mathcal{M}_m} \mathrm{P}^\mu_{s_0,\infty}(\bar{A}) \leq \epsilon\}$$

with safety level $1 - \epsilon$.

Let us aim at computing this maximally safe policy again by means of a dynamic programming scheme. In addition, as it is reasonable in an infinite horizon setting for a time-invariant system, let us try to investigate if such a policy can be selected among the stationary Markov policies.

In the following, let us focus on the interpretation based on the expression for the probability $\mathrm{P}^\mu_{\pi,\infty}(\bar{A})$ in terms of the max cost: $\mathrm{P}^\mu_{\pi,\infty}(\bar{A}) = E^\mu_{\pi,\infty}\big[\max_{k\geq 0} \mathbf{1}_{\bar{A}}(\mathbf{s}(k))\big]$.

Unlike the additive-cost, the max cost framework is not of widespread usage and has not been deeply investigated. Extending the results developed for the infinite horizon additive cost case to the infinite horizon max cost case requires some attention regarding the following aspects:

1. it is first necessary to take care of the asymptotic properties of the max cost function at the limit;

2. and the measurability properties of the limit function, its behavior when minimized, the existence and properties of the optimal argument have to be carefully assessed.

An iterative procedure to compute $\mathrm{P}^{\mu}_{\pi,\infty}(\bar{A})$ is again possible. Conditions that yield a stationary optimal Markov policy can also be provided. Before presenting the main result, let us introduce the following change of variables, for any $N > 0$: $l = N - k$. According to this, the initial condition, chosen at $l = 0$, for any policy $\mu = (\mu_0, \mu_1, \ldots)$, is $W^{\mu}_0(s) = \mathbf{1}_A(s)$, $s \in \mathcal{S}$.

Theorem 7. *Define the maps $W^{\star}_l : \mathcal{S} \to [0,1]$, $l \geq 0$, by the recursion:*

$$W^{\star}_{l+1}(s) = \inf_{(u,\sigma)\in\mathcal{U}\times\Sigma} \big(\boldsymbol{1}_{\bar{A}}(s) + \boldsymbol{1}_A(s) H(s,(u,\sigma),W^{\star}_l)\big),$$

$s \in \mathcal{S}$, initialized with $W^{\star}_0(s) = \boldsymbol{1}_{\bar{A}}(s)$, $s \in \mathcal{S}$.
Suppose that $\exists \bar{l} \geq 0$ s.t. $U_l(s,\lambda) = \{(u,\sigma) \in \mathcal{U}\times\Sigma : H(s,(u,\sigma),W^{\star}_l) \leq \lambda\}$ is compact for all $s \in A$, $\lambda \in \mathbb{R}$, $l \geq \bar{l}$. Then, $W^{\star}_{\infty}(s) = \inf_{\mu\in\mathcal{M}_m} \mathrm{P}^{\mu}_{s,\infty}(\bar{A})$, $s \in \mathcal{S}$. Furthermore, there exists a maximally safe stationary Markov policy $\mu^{\star} = (\bar{\mu}^{\star}, \bar{\mu}^{\star}, \ldots)$, with $\bar{\mu}^{\star} : \mathcal{S} \to \mathcal{U}\times\Sigma$, given by

$$\bar{\mu}^{\star}(s) = \arg \inf_{(u,\sigma)\in\mathcal{U}\times\Sigma} H(s,(u,\sigma),W^{\star}_{\infty}),\ \forall s \in A. \qquad \square$$

Proof. For any $N \geq 0$, function $W^{\star}_N$ can be shown to be universally measurable and lower semianalytic. This is easily obtained by induction, applying the same line of reasoning as in Theorem 6 for proving Assumption A2 and taking into consideration the initialization $W^{\star}_0(s) = \mathbf{1}_{\bar{A}}(s)$, $s \in \mathcal{S}$ of the iterative procedure to compute $W^{\star}_N$. Observe now that the sequence of functions $W^{\star}_N(\cdot)$, $N \geq 0$, is increasing and upper bounded: $W^{\star}_N(s) \leq W^{\star}_{N+1}(s) \leq 1$, $\forall s \in \mathcal{S}$, $\forall N \geq 0$, ($W^{\star}_N(s)$ is the minimum probability of entering the unsafe set $\bar{A}$ starting from s along the time horizon $[0,N]$). Then, the limit function $W^{\star}_{\infty}$ will also be lower semi-analytic (ref. [Bertsekas and Shreve, 1996, Lemma 7.30(2)]), hence universally measurable. By [Bertsekas and Shreve, 1996, Prop. 9.17], it can be proven that, for any $s \in \mathcal{S}$,

$$W^{\star}_{\infty}(s) = \inf_{(u,\sigma)\in\mathcal{U}\times\Sigma} \big(\mathbf{1}_{\bar{A}}(s) + \mathbf{1}_A(s) H(s,(u,\sigma),W^{\star}_{\infty})\big).$$

Exploiting [Bertsekas and Shreve, 1996, Prop. 5.8], it is possible to show that the dynamic programming scheme has a fixed point and that $W^{\star}_{\infty}(\cdot)$ is the limit function;

thus it is the infimum in the statement. Finally, by [Bertsekas and Shreve, 1996, Cor. 9.12.1], it is possible to conclude that there exists a stationary optimal policy, the form of which is that in the statement. This concludes the proof. □

Alternative, extended proof. Let us start by studying some measurability issues. Measurability and analyticity properties are directly transferred from Theorems 5 and 6 as follows: we have shown that each $W_l^{\mu}(s), \forall s \in \mathcal{S}$, and in particular $W_l^{\star}(s)$, are universally measurable and lower semi-analytic, for any $l \in [0, N]$. By [Bertsekas and Shreve, 1996, Lemma 7.30(2)], the limit $W_{\infty}^{\mu}(s)$ will be as well lower semi-analytic. Furthermore, $W_{\infty}^{\star}(s)$ will also be lower semi-analytic [Bertsekas and Shreve, 1996, Prop. 7.47]. Being a semi-analytic function analytically measurable [Bertsekas and Shreve, 1996, Sec. 7.7], it follows that $W_{\infty}^{\star}(s)$ is also endowed with universal measurability.

Let us now turn our attention to some structural considerations. Observe that, fixing any particular $s \in \mathcal{S}$ and a policy $\mu = (\mu_0, \mu_1, \ldots) \in \mathcal{M}_m$, the cost function $W_k^{\mu}(s)$ is monotonic with time. More precisely, the following holds:

$$W_0^{\mu}(s) \leq W_k^{\mu}(s) \leq W_{\infty}^{\mu}(s), \quad \forall k > 0, \ \forall s \in \mathcal{S}. \tag{2.26}$$

This increasing property also encompasses the *monotonicity assumption*, as defined in [Bertsekas and Shreve, 1996, Sec. 6.1], which is retained at each time interval thanks to the structure of the cost functions.

From the monotonically increasing property of equation (2.26), in particular, it follows that, for an optimal policy $\mu^{\star} = (\mu_0^{\star}, \mu_1^{\star}, \ldots)$, $W_k^{\star}(s) \leq W_{k+1}^{\star}(s) \leq W_{\infty}^{\star}(s), \forall k > 0$. This, in conjunction with the property that $0 \leq W_k^{\star}(s) \leq 1, \forall k > 0, \forall s \in \mathcal{S}$, ensures that the sequence $\{W_k^{\star}(s)\}_{k=0,1,\ldots}$ will converge to a fixed $W_{\infty}^{\star}(s)$. Furthermore, the following fact is particularly interesting:

$$W_{\infty}^{\star}(s) = \inf_{(u,\sigma) \in \mathcal{U} \times \Sigma} \left(\mathbf{1}_{\bar{A}}(s) + \mathbf{1}_A(s) H(s, (u, \sigma), W_{\infty}^{\star}) \right). \tag{2.27}$$

In words, the application of the dynamic programming scheme yields a convergent cost function. At this point, we need to show that the optimal cost function attained by the application of the dynamic programming scheme is in fact the one obtained

optimizing the infinite horizon cost function directly over all the possible policies $\mu \in \mathcal{M}_m$: this is obtained by use of [Bertsekas and Shreve, 1996, Proposition 9.10, section 5.4].

Next, let us show that the dynamic programming scheme yields a stationary optimum. Picking any $\lambda \in \mathbb{R}, U_k(s, \lambda)$ is compact for all $k > 0$ by assumption. Furthermore, from the monotonicity property of the cost functions it follows that, in particular, it is possible to select $\lambda = W_\infty^\star(s)$. This suggests that all the points of the sequence $\mu \in \mathcal{M}_m, \{\mu_k\}_{k=(0,1,\dots)}$ are contained in the compact set $\bigcap_{k=\{0,1,\dots\}} U_k(s, W_\infty^\star(s)), \forall s \in \mathcal{S}$. Hence, its limit point $\bar{\mu}$ is as well. In particular, (2.27) holds for $\bar{\mu} = (\bar{u}, \bar{\sigma})$. This says that, for any $s \in \mathcal{S}, k > 0$, there exists a $\bar{\mu} = (\bar{u}, \bar{\sigma}) \in \mathcal{U} \times \Sigma$ such that

$$\bar{\mu} = \mu^\star = \arg \inf_{(u,\sigma)\in\mathcal{U}\times\Sigma} H(s, (u, \sigma), W_k^\star).$$

We have explicitly shown that the infimization of the quantities $W_k^\mu(s)$ will yield a result, for all $s \in \mathcal{S}$ and $k > 0$. This is what in Theorems 5 and 6 is called the *exact selection assumption* and is proved by the Jankov-von Neumann theorem and [Bertsekas and Shreve, 1996, section 6.2 and Proposition 7.50]. In addition, the optimal control $\bar{\mu}$ will be stationary [Bertsekas and Shreve, 1996, Propositions 9.9 and 9.12]. This concludes the proof. □

2.2.5 Regulation and Practical Stabilization Problems

Reachability can also be studied within the framework of *regulation* theory where the aim is to steer the state of the system close to some desired operating condition. This can be achieved by considering a small neighborhood around the desired operating condition, and by solving a reachability problem with a time-varying region that shrinks to that neighborhood as the "safe" set for the system.

If the state of the system has then to be maintained within this neighborhood indefinitely, [Bertsekas, 1972], one can split the problem into a finite horizon time-varying reachability problem and a subsequent infinite horizon one. This approach has close connections with control design for *practical stabilization*, [Picasso and Bicchi, 2005]. The application of reachability analysis to regulation problems for DTSHS

is discussed in section 2.3.2.2 with reference to a simple application example, where the problem is to drive the temperature of a room close to some desired value by controlling a heater.

2.2.6 Other Related Control Problems

Fault-Tolerant Control. A further possible application of the dynamic programming approach to stochastic reachability is the verification of *fault-tolerant control* schemes, within the DTSHS setting. By introducing a policy that selects the worst faulty action based on the current value of the state to model the faulty system behavior, one can verify if the state of the system will remain within a well-behaved set with high probability, despite of the possible occurrence of faults. In this application context, it may be useful to enlarge the class of considered policies from deterministic to *randomized* policies since the occurrence of faults is typically governed by a probabilistic law rather than a deterministic (quasi-adversarial) behavior.

Adversarial Uncertainty. Engineering Systems can be characterized not only by random uncertainty, coming either from noise or from lack of perfect knowledge of the system or of its environment. It is often also the case that the system may include an adversarial input, which acts on the system with an objective that is either different, or in competition, with the one of the main actor. An example of such a system comes from Air Traffic Control or general military applications [Ghosh and Marcus, 1995; Lygeros *et al.*, 1997; Tomlin *et al.*, 2000; Bayen *et al.*, 2003].

Mathematically, this can be described by introducing a game-theoretical framework, rather than the known optimal control one. Some classic literature on the topic come from [Issacs, 1967; Evans and Souganidis, 1984; Lagunov, 1985; Fleming and Souganidis, 1989].

2.2.7 Embedding Performance into the Problem

More challenging extensions of the approach include the treatment of general *optimal control problems with safety constraints*, as in [Batina *et al.*, 2002]. Quite common is that the controlled system is required to behave optimally or efficiently according to

some performance criterion, while evolving within a safe or desired set. Related to this, an interesting approach is that proposed in [Lygeros *et al.*, 1999], addressing a multi-objective control problem with requirements ranking, for deterministic hybrid systems. Other related work on the subject of hierarchical or multi-objective control can be found in [Tomlin *et al.*, 1998a; Oishi *et al.*, 2001; Lygeros, 2004b].

In some cases, such as when the system performance is evaluated in terms of an additive cost, a relaxed version of the optimal constrained control problem can be formulated where a new state component representing at each time instant the cumulated cost is introduced and the objective is maintaining the system within an extended safe region where the cost is sufficiently low. As an example, suppose that in the considered temperature regulation case study it is desirable to limit the number of switchings. It is then possible to assign a unitary cost to each commutation, add a state variable counting the number of commutations, and keep this new variable within a bounded region around the origin, with an upper limit corresponding to the total number of allowed commutations. Clearly, the caveat to this approach is that an increase in dimensionality may turn out to be computationally unfeasible in practice.

2.2.8 State and Control Discretization

This section unfolds from the consideration that since an analytic solution to the dynamic programming scheme (such as that in Theorem 5, or in 6) is generally hard to find, the computational aspects to the problem are of key importance to its actual implementation. This is particularly true for the case under study, where the general dynamics and the nonlinear shape of the cost function do not enable a closed-form expression of the quantities of interest.

The above issue can be tackled in two different ways. One can resort to a numerical approximation scheme relying on the discretization of the continuous state and control input spaces. Alternatively, one can introduce a family of finitely parameterized functions, and then look for the cost-to-go function within that family according to the neuro-dynamic programming approach, as in [Bertsekas and Tsitsiklis, 1996]. Here, the first option is considered.

Let us study a discretization procedure for the numerical solution to the dynamic

programming equation of either of the stochastic reachability problems described in the previous sections. The convergence of the numerical solution to the actual solution as the grid size goes to zero will be assessed. Additionally, explicit bounds on the level of approximation introduced for a given small but nonzero grid size will be derived. The study is inspired by the reference work [Bertsekas, 1975], which discusses discretization procedures for the numerical solution to dynamic programming in the additive cost case and for stochastic– non hybrid– systems. The present contribution extends this approach to a hybrid system setting with multiplicative cost and general disturbances. The current approach is not the unique attempt within the set of approximation techniques. For instance, a numerical approximation scheme was proposed in [Prandini and Hu, 2006a] for estimating the probability of remaining within a safe set for a certain class of autonomous, continuous time SHS, namely switching diffusions. The discretization process in that case involved gridding the system both in time and in space, and is based on weak convergence approximation methods studied in [Kushner, 1984]. Asymptotic convergence of the estimate to the true probability was proven, but no bounds were provided for assessing the quality of the derived non-asymptotic estimates.

For notational simplicity, let us consider the following case:

Assumption 7. *The control space $\mathcal{A} = \mathcal{U} \times \Sigma$ is a finite set.* □

Under this assumption, the compactness condition required in Theorem 5 (or in Theorem 6) for the existence of a maximally safe Markov policy is not required. Some comments on the generalization of the results to the case when $\mathcal{A}$ is a compact uncountable set in an Euclidean space is postponed to the conclusion section.

Let us focus on the max approach in Theorem 5 and recall the DP scheme in equations (2.18)-(2.19), involved in the computations of the maximal probabilistically safe sets and of the maximally safe policies. For the purpose of numerical approximation, it is important to note that the scheme can be restricted to the safe set A, assmed to be compact, of the state space as follows:

$$V_k^*(s) = \max_{a \in \mathcal{A}} \int_A V_{k+1}^*(s_{k+1}) T_s(ds_{k+1}|s,a), \; s \in A, \tag{2.28}$$

initialized with $V_N^*(s) = 1$, $s \in A$, and

$$\mu_k^*(s) = \arg\max_{a \in \mathcal{A}} \int_A V_{k+1}^*(s_{k+1}) T_s(ds_{k+1}|s,a), \; \forall s \in A. \tag{2.29}$$

This is quite intuitive, since for values of the state outside A the cost-to-go function is identically zero for any μ. Moreover, the optimal policy $\mu^* : \mathcal{S} \to \mathcal{A}$ can be set arbitrarily on the set $\mathcal{S} \setminus A$. Thus, we just have to consider the values for the state within the compact set A. The advantage of considering only the compact set A is that we can adopt a finite discretization for the continuous state component within A in the numerical approximation scheme for solving the dynamic programming algorithm and determining the optimal policy μ^*. The idea is that, under suitable regularity conditions on the transition kernels, the optimal cost-to-go functions can be shown to be Lipschitz continuous in the continuous state component over A. This property (valid only within A, given the discontinuity when passing from a safe state within A to an unsafe state outside A) is used for determining bounds to the numerical approximated solution.

2.2.8.1 Discretization Procedure

State discretization. As discussed before, let us restrict computations–and the discretization procedure–to the compact safe set A. The set $A \subset \mathcal{S}$ is given by $A = \bigcup_{q \in \mathcal{Q}} \{q\} \times X^q$. The size of the continuous state space within A is measured by $\lambda := \max_{q \in \mathcal{Q}} \texttt{Leb}(X_q)$, where $\texttt{Leb}(X_q)$ denotes the Lebesgue measure of the set $X_q \subset \mathbb{R}^{n(q)}$. For simplicity, let us assume that the compact set X^q is $X^q \neq \varnothing$, for all $q \in \mathcal{Q}$. Let us introduce a partition of cardinality m_q of set $X^q \subset \mathbb{R}^{n(q)}$, $q \in \mathcal{Q}$: $X^q = \bigcup_{i=1}^{m_q} X_i^q$, where X_i^q, $i = 1, \ldots, m_q$ are pair-wise disjoint Borel sets $X_i^q \in \mathcal{B}(\mathbb{R}^{n(q)})$, $X_i^q \cap X_j^q = \varnothing$, $\forall i \neq j$. For any q and i, pick a hybrid state value $v_i^q \in \{q\} \times X_i^q$. The set of all these discrete values in the hybrid state is the grid $\mathcal{G} := \{v_i^q, i = 1, \ldots, m_q, \, q \in \mathcal{Q}\}$. Notice that the compactness assumption on A ensures the finiteness of the cardinality of $\mathcal{G}$. Denote with $d_{q,i}$ the diameter of the set X_i^q, $d_i^q = \sup\{\|x - x'\| : x, x' \in X_i^q\}$. Then, $\Delta_s := \max_{i=1,\ldots,m_q, q \in \mathcal{Q}} d_i^q$ is the *grid size parameter*.

Note that, unlike [Bertsekas, 1975] where the system dynamics are described

through a difference nonlinear equation affected by a stochastic disturbance, we do not have a disturbance input appearing explicitly. The definition of the dynamics of the system via stochastic kernels incorporates both the disturbance effect and the deterministic contribution to the system evolution. As a consequence, by discretizing the state space, a discretization of the disturbance space is also implicitly defined (see the example in section 2.3.3).

Dynamic programming approximation. With reference to the finite state grid $\mathcal{G}$ and the scheme in Theorem 5, equation (2.18), let us introduce a discretized version of the dynamic programming equations as follows:

$$\begin{cases} \widehat{V}_N^*(v_i^q) = 1, \text{ if } v_i^q \in \mathcal{G} \\ \widehat{V}_N^*(s) = \widehat{V}_N^*(v_i^q), \text{ if } s \in \{q\} \times X_i^q, \forall i = 1, \dots, m_q,\ q \in \mathcal{Q} \\ \widehat{V}_k^*(v_i^q) = \max_{a \in \mathcal{A}} \int_A \widehat{V}_{k+1}^*(s) T_s(ds|v_i^q, a), \quad \text{if } v_i^q \in \mathcal{G} \\ \widehat{V}_k^*(s) = \widehat{V}_k^*(v_i^q), \text{ if } s \in \{q\} \times X_i^q, i = 1, \dots, m_q,\ q \in \mathcal{Q}, \text{ for } k \in [0, N-1]. \end{cases} \tag{2.30}$$

Note that due to the piecewise constant approximation of the optimal cost-to-go function and to the definition of T_s in equation (2.1), the integral in equation in (2.30) can be rewritten as

$$\widehat{V}_k^*(v_i^q) = \max_{a \in \mathcal{A}} \Big\{ \sum_{j=1,\dots,m_q} \widehat{V}_{k+1}^*(v_j^q) T_q(q|v_i^q, a) \int_{X_j^q} T_x(dx|v_i^q, a) + \sum_{\substack{j=1,\dots,m_{\bar{q}}, \\ \bar{q} \neq q \in \mathcal{Q}}} \widehat{V}_{k+1}^*(v_j^{\bar{q}}) T_q(\bar{q}|v_i^q, a) \int_{X_j^{\bar{q}}} R(dx|v_i^q, a, \bar{q}) \Big\},$$

which explicitly shows that (2.30) consists of a computation on the finite grid $\mathcal{G}$.

Based on the approximated optimal cost-to-go $\widehat{V}^*$ let us define a piecewise constant Markov policy $\widehat{\mu}^* = (\widehat{\mu}_0^*, \dots, \widehat{\mu}_{N-1}^*)$, $\widehat{\mu}_k^* : \mathcal{S} \to \mathcal{A}$, $k \in [0, N-1]$, as follows:

$$\begin{cases} \widehat{\mu}_k^*(v_i^q) = \arg\max_{a \in \mathcal{A}} \int_A \widehat{V}_{k+1}^*(s) T_s(ds|v_i^q, a), \quad \text{if } v_i^q \in \mathcal{G}, \\ \widehat{\mu}_k^*(s) = \widehat{\mu}^*(v_i^q), \quad \text{if } s \in \{q\} \times X_i^q, i = 1, \dots, m_q,\ q \in \mathcal{Q}. \end{cases} \tag{2.31}$$

As previously argued in general for any policy, also $\widehat{\mu}^*$ can be be arbitrarily selected outside the safe set A.

The actual performance of the so-obtained policy $\widehat{\mu}^*$ are given by the corresponding values for the original cost-to-go functions $V_k^{\widehat{\mu}^*}$, $k = 0, 1, \ldots, N$ (notice the difference with the terms in 2.30), that can be computed by the recursion in Lemma 1. In particular, $V_0^{\widehat{\mu}^*}(s)$, $s \in A$, provides the value of the probability that the system will remain within A in the time horizon $[0, N]$ starting from s under policy $\widehat{\mu}^*$.

In the following section, it will be shown that under proper assumptions the performance of policy $\widehat{\mu}^*$ tends to the one of a maximally safe policy μ^*, as the grid size parameter Δ_s goes to zero.

2.2.8.2 Convergence Analysis

Let us suppose that the stochastic kernels T_x and R on the continuous component of the hybrid state in Definition 13 of the DTSHS admit density t_x and r. Let us further assume that the densities t_x and r, as well as the stochastic kernel T_q, satisfy the following Lipschitz condition.

Assumption 8. *Let the following hold:*

1. $|T_q(\bar{q}|s,a) - T_q(\bar{q}|s',a)| \leq k_1\|x - x'\|$, *for all* $s = (q,x), s' = (q,x') \in A$, $a \in \mathcal{A}$, *and* $\bar{q} \in \mathcal{Q}$,
2. $|t_x(\bar{x}|s,a) - t_x(\bar{x}|s',a)| \leq k_2\|x - x'\|$, *for all* $s = (q,x), s' = (q,x') \in A$, $a \in \mathcal{A}$, *and* $(q,\bar{x}) \in A$,
3. $|r(\bar{x}|s,a,\bar{q}) - r(\bar{x}|s',a,\bar{q})| \leq k_3\|x - x'\|$, *for all* $s = (q,x), s' = (q,x') \in A$, $a \in \mathcal{A}$, $(\bar{q},\bar{x}) \in A$, *and* $\bar{q} \neq q$,

where k_1, k_2 *and* k_3 *are suitable finite Lipschitz constants.* □

Based on this assumption, it is possible to prove that the optimal cost-to-go functions satisfy some Lipschitz condition over A. This property will be fundamental in proving the convergence result.

Theorem 8. *Under Assumption 8, the optimal cost-to-go functions satisfy the following Lipschitz condition over* A*:*

$$|V_k^*(s) - V_k^*(s')| \leq \mathcal{K}\|x - x'\|, \forall\, s = (q,x),\ s' = (q,x') \in A, \tag{2.32}$$

for any $k \in [0, N]$. *The constant* $\mathcal{K}$ *is given by* $\mathcal{K} = mk_1 + \lambda\big(k_2 + (m-1)k_3\big)$. $\square$

Proof. Since $V_N^*(s) = V_N^*(s') = 1$, for all $s, s' \in A$, then, the inequality in (2.32) is trivially satisfied for $k = N$. For any $k \in [0, N-1]$, $s = (q, x)$, $s' = (q, x') \in A$, from the dynamic programming equation (2.18) and definition (2.1) of T_s, we have:

$$
\begin{aligned}
&|V_k^*(s) - V_k^*(s')| = \qquad (2.33)\\
&= \Big| \max_{a\in\mathcal{A}} \Big\{ \int_A V_{k+1}^*(s_{k+1}) T_s(ds_{k+1}|s,a) \Big\} - \max_{a\in\mathcal{A}} \Big\{ \int_A V_{k+1}^*(s_{k+1}) T_s(ds_{k+1}|s',a) \Big\}\\
&= \Big| \max_{a\in\mathcal{A}} \Big\{ T_q(q|s,a) \int_{X^q} V_{k+1}^*(q,\bar{x}) T_x(d\bar{x}|s,a) + \sum_{\bar{q}\neq q} T_q(\bar{q}|s,a) \int_{X^{\bar{q}}} V_{k+1}^*(\bar{q},\bar{x}) R(d\bar{x}|s,a,\bar{q}) \Big\}\\
&- \max_{a\in\mathcal{A}} \Big\{ T_q(q|s',a) \int_{X^q} V_{k+1}^*(q,\bar{x}) T_x(d\bar{x}|s',a) + \sum_{\bar{q}\neq q} T_q(\bar{q}|s',a) \int_{X^{\bar{q}}} V_{k+1}^*(\bar{q},\bar{x}) R(d\bar{x}|s',a,\bar{q}) \Big\} \Big|\\
&\leq \max_{a\in\mathcal{A}} \Big\{ \Big| T_q(q|s,a) \int_{X^q} V_{k+1}^*(q,\bar{x}) T_x(d\bar{x}|s,a) - T_q(q|s',a) \int_{X^q} V_{k+1}^*(q,\bar{x}) T_x(d\bar{x}|s',a) \Big|\\
&+ \sum_{\bar{q}\neq q} \Big| T_q(\bar{q}|s,a) \int_{X^{\bar{q}}} V_{k+1}^*(\bar{q},\bar{x}) R(d\bar{x}|s,a,\bar{q}) - T_q(\bar{q}|s',a) \int_{X^{\bar{q}}} V_{k+1}^*(\bar{q},\bar{x}) R(d\bar{x}|s',a,\bar{q}) \Big| \Big\}.
\end{aligned}
$$

Let us next show two intermediate results that will be useful for proving the Lipschitz property for V_k^*. The following chain of inequalities can be easily proven using the fact that $|V_{k+1}^*(q,\bar{x})| \leq 1$ and Assumption 8:

$$
\begin{aligned}
&\Big| \int_{X^q} V_{k+1}^*(q,\bar{x}) T_x(d\bar{x}|(q,x),a) - \int_{X^q} V_{k+1}^*(q,\bar{x}) T_x(d\bar{x}|(q,x'),a) \Big|\\
&\quad \leq \Big| \int_{X^q} V_{k+1}^*(q,\bar{x}) \big(T_x(d\bar{x}|(q,x),a) - T_x(d\bar{x}|(q,x'),a)\big) \Big|\\
&\quad \leq \int_{X^q} |V_{k+1}^*(q,\bar{x})| \big| T_x(d\bar{x}|(q,x),a) - T_x(d\bar{x}|(q,x'),a) \big|\\
&\quad \leq \int_{X^q} \big| T_x(d\bar{x}|(q,x),a) - T_x(d\bar{x}|(q,x'),a) \big|\\
&\quad \leq \lambda k_2 \|x - x'\|, \forall (q,x), (q,x') \in A, a \in \mathcal{A}.
\end{aligned}
$$

Similarly, we have that

$$\Big| \int_{X^{\bar{q}}} V_{k+1}^*(\bar{q},\bar{x}) R(d\bar{x}|(q,x),a,\bar{q}) - \int_{X^{\bar{q}}} V_{k+1}^*(\bar{q},\bar{x}) R(d\bar{x}|(q,x'),a,\bar{q}) \Big| \\ \leq \lambda k_3 \|x - x'\|, \forall (q,x),(q,x') \in A, \bar{q} \in \mathcal{Q}, a \in \mathcal{A}.$$

Recall now that the product of two functions $\alpha, \beta : \mathcal{E} \to \mathbb{R}$ that are Lipschitz continuous over a compact set $\mathcal{E}$ of a Euclidean space, with Lipschitz constants respectively k_α and k_β, satisfies:

$$|\alpha(z_1)\beta(z_1) - \alpha(z_2)\beta(z_2)| \leq \left\{ k_\alpha \sup_{z\in\mathcal{E}} |\beta(z)| + k_\beta \sup_{z\in\mathcal{E}} |\alpha(z)| \right\} \|z_1 - z_2\|.$$

By this inequality, applied twice to the two terms on the rhs of equation (2.33) with $\alpha(z) = T_q(q|(q,z),a)$ both times, and with either $\beta(z) = \int_{X^q} V_{k+1}^*(q,\bar{x}) T_x(d\bar{x}|(q,z),a)$ or $\beta(z) = \int_{X^{\bar{q}}} V_{k+1}^*(q,\bar{x}) R(d\bar{x}|(\bar{q},z),a,\bar{q})$, it is easily shown that

$$|V_k^*(s) - V_k^*(s')| \leq \big[m k_1 + \lambda\big(k_2 + (m-1)k_3\big)\big] \|x - x'\|,$$

which concludes the proof. □

Finally, let us prove the following convergence result, relating the error on the optimal cost functions, as well as the error on the use of the approximate optimal policy in the optimal cost function.

Theorem 9. *Under Assumption 8, there exists positive constants γ_k, $k = 0,\dots,N$, such that the solutions $\widehat{V}_k^*$ to the approximated dynamic programming equations (2.30) and the cost-to-go functions of the corresponding Markov policy $\widehat{\mu}^*$ defined in (2.31) satisfy:*

$$\big|V_k^*(s) - \widehat{V}_k^*(s)\big| \leq \gamma_k \Delta_s, s \in A, \tag{2.34}$$

$$\big|V_k^*(s) - V_k^{\widehat{\mu}^*}(s)\big| \leq \nu_k \Delta_s, s \in A. \tag{2.35}$$

where $\gamma_k = \gamma_{k+1} + \mathcal{K}$, $k = 1,2\dots,N-1$, initialized with $\gamma_N = 0$, and $\nu_k = \gamma_k + \gamma_{k+1} + \mathcal{K} + \nu_{k+1}$, $k = 1,2\dots,N-1$, initialized with $\nu_N = 0$, and where $\mathcal{K} = mk_1 +$

$\lambda\big(k_2 + (m-1)k_3\big)$. □

Proof. Let us prove equation (2.34) by induction. Since $V_N^*(s) = \widehat{V}_N^*(s) = 1$, $s \in A$, then, equation (2.34) trivially holds for $k = N$.
Let us suppose by induction hypothesis that $\big|V_{k+1}^*(s) - \widehat{V}_{k+1}^*(s)\big| \leq \gamma_{k+1}\Delta_s$, $s \in A$. Then, for any $v_i^q \in \{q\} \times X_i^q \subset A$, $i = 1, \ldots, m_q$, $q \in \mathcal{Q}$, we have:

$$
\begin{aligned}
&|V_k^*(v_i^q) - \widehat{V}_k^*(v_i^q)| = \\
&\Big|\max_{a\in\mathcal{A}}\Big\{\int_A V_{k+1}^*(s_{k+1})T_s(ds_{k+1}|v_i^q,a)\Big\} - \max_{a\in\mathcal{A}}\Big\{\int_A \widehat{V}_{k+1}^*(s_{k+1})T_s(ds_{k+1}|v_i^q,a)\Big\}\Big| \\
&\leq \max_{a\in\mathcal{A}}\Big|\int_A V_{k+1}^*(s_{k+1})T_s(ds_{k+1}|v_i^q,a) - \int_A \widehat{V}_{k+1}^*(s_{k+1})T_s(ds_{k+1}|v_i^q,a)\Big| \\
&\leq \max_{a\in\mathcal{A}}\int_A \big|V_{k+1}^*(s_{k+1}) - \widehat{V}_{k+1}^*(s_{k+1})\big|T_s(ds_{k+1}|v_i^q,a) \leq \gamma_{k+1}\Delta_s.
\end{aligned}
$$

Thus, for any $s \in \{q\} \times X_i^q \subset A$, $i = 1, \ldots, m_q$, $q \in \mathcal{Q}$:

$$
\begin{aligned}
|V_k^*(s) - \widehat{V}_k^*(s)| = |V_k^*(s) - \widehat{V}_k^*(v_i^q)| &\leq |V_k^*(s) - V_k^*(v_i^q)| + |V_k^*(v_i^q) - \widehat{V}_{N-1}^*(v_i^q)| \\
&\leq \mathcal{K}\Delta_s + \gamma_{k+1}\Delta_s = \gamma_k\Delta_s,
\end{aligned}
$$

where Theorem 8 has been used. This concludes the proof on equation (2.34).

As for equation (2.35), let us start observing that it trivially holds for $k = N$ since $V_N^*(s) = V_N^{\widehat{\mu}^*}(s) = 1$, $s \in A$. Note that, for any $s \in \{q\} \times X_i^q \subset A$, $i = 1, \ldots, m_q$, $q \in \mathcal{Q}$,

$$
\begin{aligned}
\big|V_k^*(s) - V_k^{\widehat{\mu}^*}(s)\big| &\leq \big|V_k^*(s) - \widehat{V}_k^*(s)\big| + \big|\widehat{V}_k^*(s) - V_k^{\widehat{\mu}^*}(s)\big| \\
&\leq \gamma_k\Delta_s + \big|\widehat{V}_k^*(v_i^q) - V_k^{\widehat{\mu}^*}(s)\big|,
\end{aligned}
$$

where the last inequality follows from equation (2.34) and the definition of $\widehat{V}_k^*$, which is constant and equal to $V_k^*(v_i^q)$ within each set $\{q\} \times X_i^q$. In view of this inequality we need to prove that

$$
\big|\widehat{V}_k^*(v_i^q) - V_k^{\widehat{\mu}^*}(s)\big| \leq (\gamma_{k+1} + \nu_{k+1} + \mathcal{K})\Delta_s. \tag{2.36}
$$

Suppose by induction hypothesis that $\big|V_{k+1}^*(s) - V_{k+1}^{\widehat{\mu}^*}(s)\big| \leq \nu_{k+1}\Delta_s, s \in A$. By

Lemma 1 we get

$$
\begin{aligned}
&\left|\widehat{V}_k^*(v_i^q) - V_k^{\widehat{\mu}^*}(s)\right| \\
&= \left| \max_{a\in\mathcal{A}} \left\{ \int_A \widehat{V}_{k+1}^*(s_{k+1}) T_s(ds_{k+1}|v_i^q, a) \right\} - \int_A \widehat{V}_{k+1}^{\widehat{\mu}^*}(s_{k+1}) T_s(ds_{k+1}|s, \widehat{\mu}^*(v_i^q)) \right| \\
&= \left| \int_A \widehat{V}_{k+1}^*(s_{k+1}) T_s(ds_{k+1}|v_i^q, \widehat{\mu}^*(v_i^q)) - \int_A \widehat{V}_{k+1}^{\widehat{\mu}^*}(s_{k+1}) T_s(ds_{k+1}|s, \widehat{\mu}^*(v_i^q)) \right| \\
&\leq \left| \int_A \widehat{V}_{k+1}^*(s_{k+1}) T_s(ds_{k+1}|v_i^q, \widehat{\mu}^*(v_i^q)) - \int_A V_{k+1}^*(s_{k+1}) T_s(ds_{k+1}|v_i^q, \widehat{\mu}^*(v_i^q)) \right| \\
&+ \left| \int_A V_{k+1}^*(s_{k+1}) T_s(ds_{k+1}|v_i^q, \widehat{\mu}^*(v_i^q)) - \int_A V_{k+1}^*(s_{k+1}) T_s(ds_{k+1}|s, \widehat{\mu}^*(v_i^q)) \right| \\
&+ \left| \int_A V_{k+1}^*(s_{k+1}) T_s(ds_{k+1}|s, \widehat{\mu}^*(v_i^q)) - \int_A \widehat{V}_{k+1}^{\widehat{\mu}^*}(s_{k+1}) T_s(ds_{k+1}|s, \widehat{\mu}^*(v_i^q)) \right|.
\end{aligned}
$$

Now, the first term in the right-hand-side of this inequality can be upper bounded as follows

$$
\begin{aligned}
&\left| \int_A \widehat{V}_{k+1}^*(s_{k+1}) T_s(ds_{k+1}|v_i^q, \widehat{\mu}^*(v_i^q)) - \int_A V_{k+1}^*(s_{k+1}) T_s(ds_{k+1}|v_i^q, \widehat{\mu}^*(v_i^q)) \right| \\
&\leq \int_A \left|\widehat{V}_{k+1}^*(s_{k+1}) - V_{k+1}^*(s_{k+1})\right| T_s(ds_{k+1}|v_i^q, \widehat{\mu}^*(v_i^q)) \leq \gamma_{k+1}\Delta_s,
\end{aligned}
$$

by using equation (2.34). As for the second term, by analogous steps than in the proof of Theorem 8, one can derive that

$$
\left| \int_A V_{k+1}^*(s_{k+1}) T_s(ds_{k+1}|v_i^q, \widehat{\mu}^*(v_i^q)) - \int_A V_{k+1}^*(s_{k+1}) T_s(ds_{k+1}|s, \widehat{\mu}^*(v_i^q)) \right| \leq \mathcal{K}\Delta_s.
$$

Finally, by using the induction hypothesis, the third term can be bounded as follows

$$
\begin{aligned}
&\left| \int_A V_{k+1}^*(s_{k+1}) T_s(ds_{k+1}|s, \widehat{\mu}^*(v_i^q)) - \int_A \widehat{V}_{k+1}^{\widehat{\mu}^*}(s_{k+1}) T_s(ds_{k+1}|s, \widehat{\mu}^*(v_i^q)) \right| \\
&\leq \int_A \left|V_{k+1}^*(s_{k+1}) - \widehat{V}_{k+1}^{\widehat{\mu}^*}(s_{k+1})\right| T_s(ds_{k+1}|s, \widehat{\mu}^*(v_i^q)) \leq \nu_{k+1}\Delta_s.
\end{aligned}
$$

These bounds complete the proof of equation (2.36). □

From this theorem it follows that the quality of the approximation by the nu-

merical procedure described in equations (2.30) and (2.31) improves as the grid size parameter Δ_s decreases. The rate of convergence is linear in Δ_s with a constant that depends on the Lipschitz constants k_1, k_2, and k_3 in Assumption 8 through the $\mathcal{K}$ constant defined in Theorem 8. This is not surprising because we are using a piecewise constant approximation of the optimal cost-to-go function and we expect that the optimal cost-to-go function is more regular as k_1, k_2, and k_3 are smaller. As the time horizon grows, the approximation error propagates. This is taken into account by the constants γ_k and ν_k in Theorem 9, which grow linearly as k decreases from N to 0, where N is the time-horizon length.

In section 2.3.3, a computational case study is developed, which aims at showing the convergence properties of the discretization scheme here introduced.

2.2.9 Connections with the Literature

This subsection concludes the theoretical investigations described in this part of the dissertation by putting in perspective the developed techniques and by bridging between the proposed methodologies and other efforts in the literature. Please refer to Figure 2.4 for a visualization of these connections.

As already discussed in the opening of section 2, the literature on reachability analysis for Stochastic Hybrid Systems mostly polarizes into two classes of SHS models: the GSHS (section 1.2.1, Figure 2.4 on the top-left corner), and the switching diffusion [Ghosh *et al.*, 1992]. This can be thought of as a subset of (Figure 2.4, top-left set).

For the GHSH model in the uncontrolled case, the concept of probabilistic reachability is formulated in a number of possible different ways (Figure 2.4, center-left): via the concept of hitting time [Krystul and Bagchi, 2004]; that of occupation measures [Bujorianu and Lygeros, 2003]; or through the use of Dirichlet forms [Bujorianu, 2004]. These studies, while interesting from a theoretical perspective, have thus far not been followed by an assessment of their computational aspects. In the controlled GSHS case, [Mitchell and Templeton, 2005] defines probabilistic reachability as a solution of a Hamilton-Jacobi PDE, which is set up in a game-theoretical framework. Successively, the problem is computationally solved by the application of level sets methods

[Mitchell, 2004] (Figure 2.4, bottom-left arrow). A related approach is taken, in the single-control case, by [Koutsoukos and Riley, 2006], where the problem is solved with a discretization technique inspired by [Kushner, 1984].

The second SHS model that has been used as a framework for the reachability analysis problem is that of switching diffusions. The work in [Hu *et al.*, 2003; Hu *et al.*, 2005; Prandini and Hu, 2006a; Prandini and Hu, 2006b] addresses the reachability problem from the viewpoint of air traffic control studies. Computationally, a Markov chain approximation is introduced to compute the probability of entering some assigned set, and weak convergence of the (space and time) discretization method is proved, based on earler work by [Kushner, 1984; Kushner and Dupuis, 2001; Kushner, 2002] (Figure 2.4, bottom-left arrow).

Other efforts in the literature approach the problem from a more limited perspective: for instance, from a controlled Markov chain model [Katoen, 2006]; or from a probabilistic dynamical model [Digailova and Kurzhanski, 2004]. These instances can be thought of as being subset classes of (Figure 2.4, top-left set). In the latter case, computational approaches to the probabilistic reachability problem are also proposed in [Lygeros and Watkins, 2003; Prajna *et al.*, 2004] (Figure 2.4, bottom-left area).

The present study starts by considering the controlled GSHS model in section 1.2.1 (Figure 2.4, top-left). Then, a weak approximation technique is introduced, which allows one to translate the original GSHS into a SHS with no spatially-triggered events, but simpler random arrivals of an inhomogeneous Poisson process (Figure 2.4, center-top arrow). This new model carries a number of benefits, among which the opportunity to express its dynamics via a set of relations that depend on properly defined random measures. In turn, this new expression yields itself to be passed through a time-discretization scheme (again, Figure 2.4, center-top arrow). The new, discrete-time solution of the obtained model has been shown to verify some weak convergence to the old solution, as the discretization step tends to zero. It is from this DTSHS model (Figure 2.4, top-right area), enhanced with controls, that the main part of the dissertation unfolds. First, the concept of probabilistic reachability is formally defined (Figure 2.4, center-right circle). Then, an optimal control problem is formulated. The associated computation technique, based on dynamic programming, is analyzed (Figure 2.4, bottom-right area). The DP problem, which necessarily has

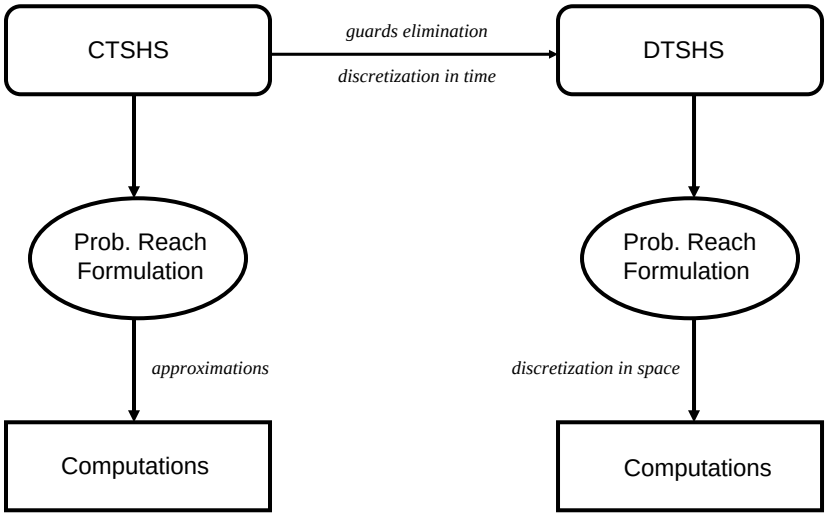

Figure 2.4: A perspective of the study.

to be computed on a discretized version of the state, control and disturbance spaces (Figure 2.4, bottom-right arrow), is shown to converge to the continuous counterpart, with proven explicit bounds.

With reference to the schematics proposed in Figure 2.4, it then makes sense to establish a comparison between the different approaches in terms of

a) the formulation of the notion of probabilistic reachability;

b) the complexity of their associated computations.

With reference to the first point, we stress that the proposed methodology is comparable, in general terms, to the techniques in [Bujorianu and Lygeros, 2003; Bujorianu, 2004; Mitchell and Templeton, 2005; Koutsoukos and Riley, 2006]. These other approaches benefit from working with continuous-time models, but this forces some restrictions on the allowed dynamics. Unlike these adjacent results, the possible time-dependence of the target set may be cleverly exploited in the present formulation to relate the current problem to other important concepts in the control literature (see section 2.2.5 for further details). The approach developed in [Hu *et al.*, 2003; Hu *et al.*, 2005; Prandini and Hu, 2006a; Prandini and Hu, 2006b], while beautifully

formalized, has the limitation of being autonomous, and with no associated resets. Furthermore, the application of weak convergence methods does not yield explicit bounds on the speed of convergence. Notice that the approximation method suffers from state space explosion issues that are relatable to the computational complexity of the techniques proposed in this thesis. This problem is further elaborated in the remainder of the paragraph.

With regards to point *b)*, a comparison is feasible with the techniques proposed in [Mitchell and Templeton, 2005; Koutsoukos and Riley, 2006]. It is fair to say that all of the approaches suffer from the curse of dimensionality that comes from the solution of DP. This poses a limitation on the practical model dimensionality that can be handled. It looks like this limit is intrinsic to the study. The model abstraction methodologies proposed in the literature do not appear to handle the controlled case in a meaningful way. On a side note, the space/time discretization techniques proposed in the related literature (see, for instance, [Mitchell, 2004]) are less prone to have convergent behavior to the actual solution, unless strict continuity requirements on the dynamics and the cost functions are raised. As it will be argued in section 2.3.4, the computational framework proposed in this thesis is instead adapted to be enhanced with some ideas which could potentially, at least in part, mitigate the computational burden which unavoidably affects the proposed methodology.

The guard elimination method introduced in section 1.3.1, and its convergence study (section 1.3.2) are original . The time-discretization technique here introduced may be compared with that in [Krystul and Bagchi, 2004], where for an autonomous switched diffusion setup a strong approximation scheme is introduced.

2.3 Computations: A Benchmark Case Study

Section 2.2 has focused on the definition of the related concepts of probabilistic reachability and safety, the understanding of these notions, and on the optimization problem that can be solved to maximize a certain criterion and to synthesize a specific controller that would yield that optimum.

In this section, the computational aspects of the solution for the above optimal control problem are investigated and discussed. This is done by considering a benchmark case study which recently appeared in the hybrid systems verification and control literature [Fehnker and Ivančić, 2004].

Although it shall be argued that, similar to other efforts which appeared in the literature, the current problem is intrinsically not "computationally friendly," (especially in its control synthesis aspects) a number of approaches will be suggested, which are aimed at in part mitigating the computational burden that the study carries along.

2.3.1 Case Study: Temperature Regulation - Modeling

Let us consider the problem of regulating the temperature in r rooms. Let us suppose that each room can be warmed by a single heater and that at most one heater can be active at a time. The problem consists in designing a control switching strategy that decides which room should be heated, based on the measurements of the temperatures of r rooms, so as to maintain the temperature of each room within a prescribed interval.

Let us next describe the controlled system through a DTSHS model $\mathscr{S}$. The precise formulation of the temperature control synthesis problem is postponed to section 2.3.2.

The system configuration is naturally described by a hybrid state, whose discrete component represents which of the r rooms is being heated, and whose continuous component represents the average temperature in each of the r rooms. The discrete state space can then be defined as $\mathcal{Q} = \{\texttt{ON}_1, \texttt{ON}_2, \ldots, \texttt{ON}_r, \texttt{OFF}\}$, where in mode $\texttt{ON}_i$ it is room i to be heated and in mode $\texttt{OFF}$ no room is heated. The map $n : \mathcal{Q} \to \mathbb{N}$,

defining the dimension of the continuous component of the hybrid state space, is the constant map $n(q) = r, \forall q \in \mathcal{Q}$.

The only available control consists in a transition input dictating which room is to be heated. Thus, the reset control space is $\Sigma = \varnothing$, whereas the transition control space is $\mathcal{U} = \{\mathtt{SW}_1, \mathtt{SW}_2, \ldots, \mathtt{SW}_r, \mathtt{SW}_{\mathtt{OFF}}\}$, where $\mathtt{SW}_i$ and $\mathtt{SW}_{\mathtt{OFF}}$ corresponds to the command of heating room i and heating no room, respectively. Note that if the system is operating in mode $\mathtt{ON}_i$ and the transition control input $\mathtt{SW}_i$ is applied, it means that the heater in room i should stay active, thus leaving the current situation unchanged. Likewise for the control $\mathtt{SW}_{\mathtt{OFF}}$ within mode $\mathtt{OFF}$.

Regarding the dynamics of the continuous state $\mathbf{x} = (\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_r)$, let us model the evolution of the average temperature $\mathbf{x}_i$ in room i by the following linear stochastic differential equation:

$$d\mathbf{x}_i(t) = \sum_{j \neq i} \tilde{a}_{ij}\big(\mathbf{x}_j(t) - \mathbf{x}_i(t)\big)dt + \tilde{b}_i\big(x_a - \mathbf{x}_i(t)\big)dt + \tilde{c}_i h_i(t)dt + d\mathbf{w}_i(t). \tag{2.37}$$

These dynamics are based to those in [Malhame and Chong, 1985], where they are however restricted to the case of a single room, controlled by a single heater. The meaning of the coefficients $\tilde{a}_{ij} = \frac{a_{ij}}{\Delta}, \tilde{b}_i = \frac{b_i}{\Delta}, \tilde{c}_i = \frac{c_i}{\Delta}$ and of the binary control function $h_i(\cdot)$ will be explained in a few lines. The term $\mathbf{w}_i(t)$ is a standard Wiener process modeling the noise affecting the temperature evolution. By applying the constant-step Euler-Maruyama scheme with constant discretization step Δ,[2] [Milstein, 1994] to (2.37), we obtain dynamic relations characterized by the following stochastic difference equation:

$$\mathbf{x}_i(k+1) = \mathbf{x}_i(k) + \sum_{j \neq i} a_{ij}\big(\mathbf{x}_j(k) - \mathbf{x}_i(k)\big) + b_i\big(x_a - \mathbf{x}_i(k)\big) + c_i h_i(k) + \mathbf{n}_i(k). \tag{2.38}$$

The meaning of the terms appearing in equation (2.38) is the following: x_a is the ambient temperature, which is assumed to be constant; b_i, a_{ij}, and c_i are non negative constants representing the average heat loss rates of room i to the ambient (b_i) and to room $j \neq i$ (a_{ij}), and the rate of heat supplied by the heater in room i (c_i),

[2] Notice that this discretization method is the one used also in section 1.4.3.

all normalized with respect to the average thermal capacity of room i and rescaled by Δ. The values taken by the a_{ij} constants reflect the rooms layout, for instance $a_{ij} = 0$ if rooms i and j are not adjacent. The term $h_i(k)$ is a Boolean function equal to 1 if $\mathbf{q}(k) = \texttt{ON}_i$ (i.e. if it is room i to be heated at time k), and equal to 0 otherwise. Furthermore, the disturbance $\{\mathbf{n}_i(k), k = 0, \ldots, N\}$ affecting the temperature evolution is a sequence of i.i.d Gaussian random variables with zero mean and variance ν^2 proportional to Δ. Let us suppose for simplicity that the disturbances affecting the temperatures of different rooms are independent: for $i \neq j$, $E[\mathbf{n}_i \mathbf{n}_j^T] = \mathbf{0}$.

Let $\mathcal{N}(\cdot; m, V)$ denote the probability measure over $(\mathbb{R}^r, \mathcal{B}(\mathbb{R}^r))$ associated with a Gaussian density function with mean m and covariance matrix V. Then, the continuous transition kernel T_x (implicitly defined via the stochastic difference equation (2.38)) can be expressed as follows:

$$T_x(\cdot\,|(q, x), u) = \mathcal{N}(\cdot; x + \Xi x + \Gamma(q), \nu^2 I), \tag{2.39}$$

where Ξ is a square matrix of size r, $\Gamma(q)$ is an r-dimensional column vector that depends on $q \in \mathcal{Q}$, and I is the identity matrix of size r.

The element in row i and column j of matrix Ξ is given by $[\Xi]_{ij} = a_{ij}$, if $j \neq i$, and $[\Xi]_{ij} = -b_i - \sum_{k \neq i} a_{ik}$, if $j = i$, for any $i = 1, \ldots, r$. As for the vector $\Gamma(q)$, its i^{th} component is $[\Gamma(q)]_i = b_i x_a + c_i$, if $q = \texttt{ON}_i$, and $[\Gamma(q)]_i = b_i x_a$, if $q \in \mathcal{Q} \setminus \{\texttt{ON}_i\}$, for any $i = 1, \ldots, r$.

Note that the transition kernel T_x governing the temperature evolution within a mode does not depend on the value u taken by the transition control input. This follows from the fact that the transition control input does not affect the temperature dynamics described in (2.38). Let us then use the notation $T_x(\cdot|(q, x))$ in place of $T_x(\cdot|(q, x), u)$.

Let us assume that during the time step when a discrete transition occurs, say from mode $\texttt{ON}_i$ to mode $\texttt{ON}_j$, the temperature keeps evolving according to the dynamics characterizing the starting condition $\texttt{ON}_i$. This is modeled by defining the reset kernel as $R(\cdot\,|(q, x), q') = T_x(\cdot\,|(q, x))$, $q, q' \in \mathcal{Q}$, $x \in \mathbb{R}^r$. Note that R does not depend on the value σ of the reset control input, since in this example the reset control space is

defined to be empty.

The transition control input affects the discrete state evolution through the discrete transition kernel T_q. In this example, discrete transitions are not influenced by the value taken by the continuous state component, so that we can take $T_q : \mathcal{Q} \times \mathcal{Q} \times \mathcal{U} \to [0,1]$, with $T_q(q'|q,u)$ representing the probability that mode q' is the successor of mode q when the transition control input u is applied. For ease of notation let us set

$$T_q(q'|q,u) = \alpha_{qq'}(u), \ q, q' \in \mathcal{Q}. \tag{2.40}$$

Thus, the discrete state evolves according to a (finite state and finite input) controlled Markov chain, which can be represented by a directed graph with $r+1$ nodes (the $r+1$ discrete modes) and $(r+1)^2$ arches, with the arch from node q to node q' labeled by the corresponding controlled transition probability $\alpha_{qq'}(u)$. Figure 2.5 represents the graph of the controlled Markov chain in the case of $r = 2$ rooms.

Not all transitions may actually occur from a node q. For instance, if the control input value $u = \mathtt{SW}_i$ is applied at node $q = \mathtt{ON}_i$, then with probability one the successor node is $q' = \mathtt{ON}_i$, because room i is currently heated and the command of heating room i is issued. If everything worked perfectly, then, the control input $u = \mathtt{SW}_i$ would lead to node $q' = \mathtt{ON}_i$ from any node q, and, similarly, $u = \mathtt{SW}_{\mathtt{OFF}}$ would lead to $q' = \mathtt{OFF}$ from any node q. The definition of the controlled transition probabilities $\{\alpha_{qq'}(u),\ q, q' \in \mathcal{Q}\}$ associated with the different $u \in \mathcal{U}$ offers the possibility of encoding delays or faulty behaviors in the commutations and structural constraints imposing, for instance, that the heat can be conveyed only from a room to a contiguous one.

2.3.2 Case Study: Temperature Regulation - Control Synthesis

2.3.2.1 Single Room Case

In this case, $r = 1$. The continuous part of the state space is one-dimensional. Consider the relations in (2.37)-(2.38). The cross-parameters a_{ij} are equal to zero, and let us name $b_1 = \frac{a}{C}\Delta, c_1 = \frac{r}{C}\Delta$ (this notation for the single room case strictly

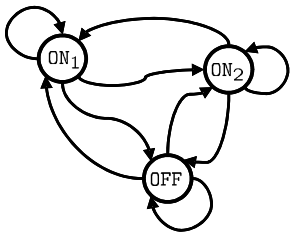

Figure 2.5: Structure of the controlled Markov chain governing the mode transitions in the case of $r = 2$ rooms.

adheres to [Malhame and Chong, 1985] by reintroducing the room average thermal capacity C and the time step Δ). The time-discretized dynamics get simplified as follows:

$$\mathbf{x}(k+1) = \begin{cases} \mathbf{x}(k) - \frac{a}{C}(\mathbf{x}(k) - x_a)\Delta + \mathbf{n}(k), & \text{if the heater is off} \\ \mathbf{x}(k) - \frac{a}{C}(\mathbf{x}(k) - x_a)\Delta + \frac{r}{C}\Delta + \mathbf{n}(k) & \text{if the heater is on.} \end{cases} \tag{2.41}$$

This dynamical relation can be equivalently expressed with the following stochastic kernel:

$$T_x(\cdot\,|(q,x)) = \begin{cases} \mathcal{N}(\cdot; x - \frac{a}{C}(x - x_a)\Delta, \nu^2), & q = \texttt{OFF} \\ \mathcal{N}(\cdot; x - \frac{a}{C}(x - x_a)\Delta + \frac{r}{C}\Delta, \nu^2), & q = \texttt{ON}. \end{cases} \tag{2.42}$$

For the sake of clarity (notice the slight semantic difference with the general definition in 2.3.1), let us name the continuous control space as $\mathcal{U} = \{0, 1\}$ with the understanding that "1" means that a switching command is issued, "0" that no switching command is issued.

Let us again assume that it takes some (random) time for the heater to actually switch between its two operating conditions, after a switching command has been issued. This is modeled by defining the discrete transition kernel T_q (and its parameters

α_{ij}) as follows

$$T_q(q'|(q,x),0) = \begin{cases} 1, & q' = q \\ 0, & q' \neq q \end{cases}$$
$$T_q(q'|(q,x),1) = \begin{cases} \alpha, & q' = \texttt{OFF},\ q = \texttt{ON} \\ 1-\alpha, & q' = q = \texttt{ON} \\ \beta, & q' = \texttt{ON},\ q = \texttt{OFF} \\ 1-\beta, & q' = q = \texttt{OFF}, \end{cases} \tag{2.43}$$

$\forall x \in \mathbb{R}$, where $\alpha \in [0,1]$ represents the probability of switching from the ON to the OFF mode in one time-step. Similarly for $\beta \in [0,1]$.

Finally, the reset kernel is modeled as previously: let us assume that the actual switching between the two operating conditions of the heater takes a time step. During this time step the temperature keeps evolving according to the dynamics referring to the starting condition.

Naïve Switching Law. Let $\bar{x}^-$, $\bar{x}^+ \in \mathbb{R}$, with $\bar{x}^- < \bar{x}^+$. Consider the (stationary) Markov policy $\mu_k : \mathcal{S} \to \mathcal{U}$ defined by

$$\mu_k((q,x)) = \begin{cases} 1, & q = \texttt{ON},\ x \geq \bar{x}^+ \text{ or } q = \texttt{OFF},\ x \leq \bar{x}^- \\ 0, & q = \texttt{ON},\ x < \bar{x}^+ \text{ or } q = \texttt{OFF},\ x > \bar{x}^- \end{cases}$$

that switches the heater on when the temperature drops below $\bar{x}^-$ and off when the temperature goes beyond $\bar{x}^+$.

Suppose that initially the heater is off and the temperature is uniformly distributed in the interval between $\bar{x}^-$ and $\bar{x}^+$, independently of the noise process affecting its evolution. In Figure 2.6, let us report some sample paths of the execution of the DTSHS associated with this policy and initial condition. The implementation is done in MATLAB. Let us plot only the continuous state realizations. The temperature is measured in Fahrenheit degrees ($^\circ F$) and the time in minutes (min). The time horizon N is taken to be 600 min. The discretization time step Δ is chosen to be 1 min.

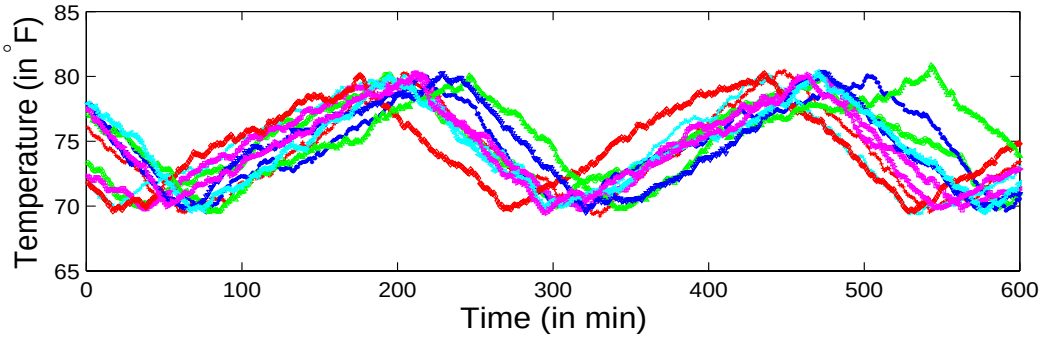

Figure 2.6: Sample paths of the temperature for the execution corresponding to a Markov policy switching the heater on/off when the temperature drops below $70°F$/goes above $80°F$, starting with heater off and temperature uniformly distributed on $[70, 80]°F$.

The parameters in equation (2.39) are assigned the following values: $x_a = 10.5°F$, $a/C = 0.1\ min^{-1}$, $r/C = 10°F/min$, and $\nu = 1°F$. The switching probabilities α and β in equation (2.43) are both chosen to be equal to 0.8. Finally, $\bar{x}^-$ and $\bar{x}^+$ are set equal to $70°F$ and $80°F$, respectively.

Note that some of the sample paths exit the set $[70, 80]°F$. This is due partly to the delay in turning the heater on/off and partly to the noise entering the system. If the objective is keeping the temperature within the set $[70, 80]$°F, more effective control policies can be found, as described in the ensuing section.

Maximally Safe Switching Law. By interpreting the desired temperatures interval as a *safe* set, we intuitively understand how this case study can be formulated as a stochastic reachability analysis problem.

More precisely, in the following three safe sets are considered: $A_1 = (70, 80)°F$, $A_2 = (72, 78)°F$, and $A_3 = (74, 76)°F$. The dynamic programming recursion described in section 2.2.3 is used to compute maximally safe policies and maximal probabilistic safe sets. The implementation is again performed in MATLAB. The temperature coordinate is discretized into 100 equally spaced values within the safe set (see section 2.2.8.1).

Figure 2.7 shows the plots of 100 temperature sample paths resulting from sampling the initial temperature from the uniform distribution over the safe sets, and using the corresponding maximally safe policy. The initial operating mode is chosen

at random between the equiprobable ON and OFF values.

It can be observed from each of the plots that the maximally safe policy computed by the dynamic programming recursion leads to an optimal behavior in the following sense: regardless of the initial state, most of the temperature sample paths tend toward the middle of the corresponding safe set. As for the A_1 and A_2 safe sets, the temperature actually remain confined within the safe set in almost all the sample paths, whereas this is not the case for A_3. The set A_3 is too small to enable the control input to counteract the drifts and the randomness in the execution in order to maintain the temperature within the safe set. The maximal probability of remaining in the safe set $p_\pi^{\mu^*}(A_i)$ for π uniform over $\mathcal{Q} \times A_i$, $i = 1, 2, 3$, is computed. The value is 0.991 for A_1, 0.978 for A_2 and 0.802 for A_3.

The maximal probabilistic safe sets $S^\star(\epsilon)$ corresponding to different safety levels $1-\epsilon$ are also calculated. The results obtained are reported in Figure 2.8 with reference to the heater initially off (plot on the left) and on (plot on the right). In all cases, as expected, the maximal probabilistic safe sets get smaller as the required safety level $1 - \epsilon$ grows. When the safe set is A_3, there is no policy that can guarantee a safety probability greater than about 0.86.

The maximally safe policies at some time instances $k \in [0, 600]$ $\mu_k^* : \mathcal{S} \to \mathcal{U}$ are shown in Figure 2.9, as a function of the continuous state and discrete state (the red crossed line refers to the OFF mode, whereas the blue circled line refers to the ON mode). The obtained result is quite intuitive. For example, at time $k = 599$, close to the end of the time horizon, and in the OFF mode, the maximally safe policy prescribes to stay in same mode for most of the continuous state values except near the lower boundary of the safe set, in which case it prescribes to change the mode to ON since there is a possibility of entering the unsafe set in the residual one-step time horizon. However, at earlier times (for instance, time $k = 1$), the maximally safe policy prescribes to change the mode even for states that are distant from the safe set boundary. Similar comments apply to the ON mode. This shows that a maximally safe policy is not stationary. By observing from top to bottom each column of Figure 2.9, one can see that this non-stationary behavior appears limited to a time interval at the end of the time horizon. This has suggested to pursue the investigation of the infinite horizon case (see section 2.2.4). Furthermore, by comparing the columns of

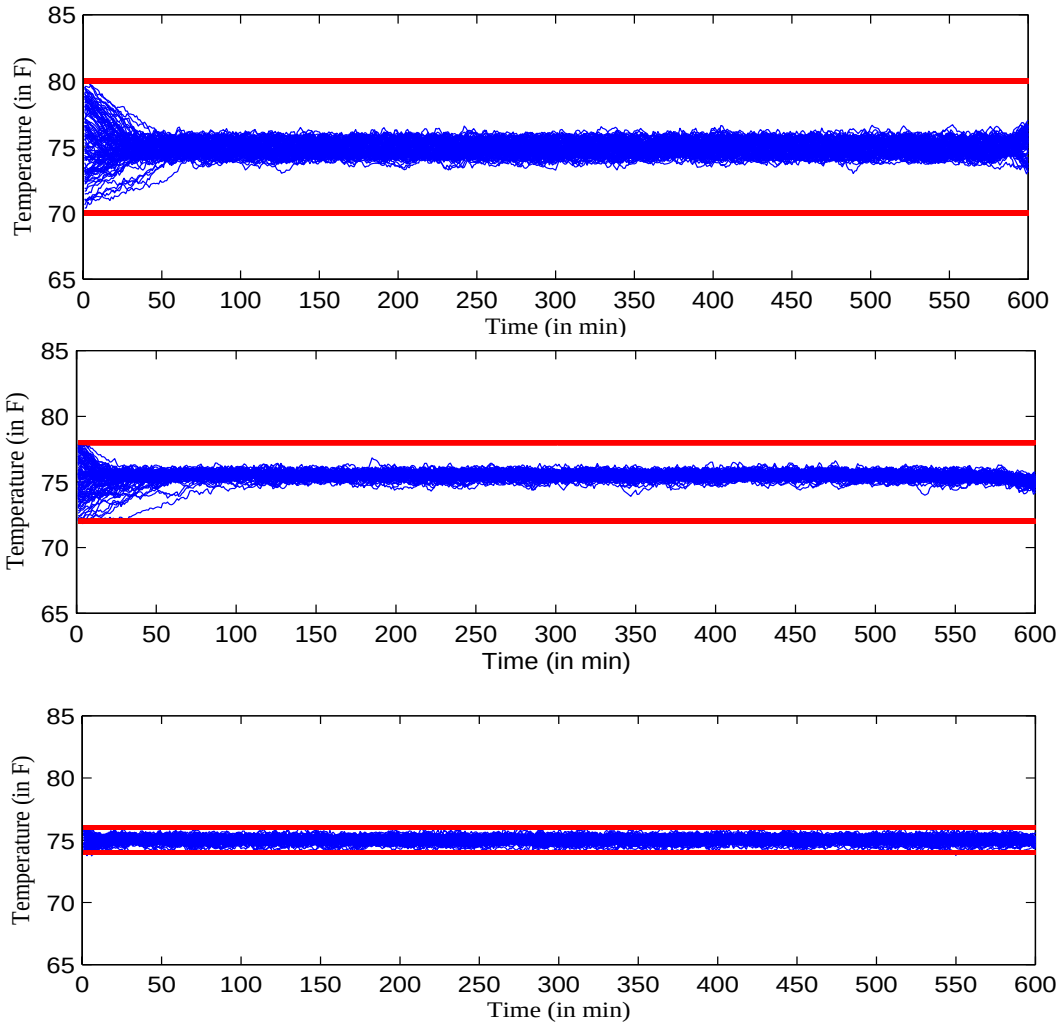

Figure 2.7: Sample paths of the temperature for the execution corresponding to maximally safe policies, when the safe set is: A_1 (top), A_2 (middle), and A_3 (bottom).

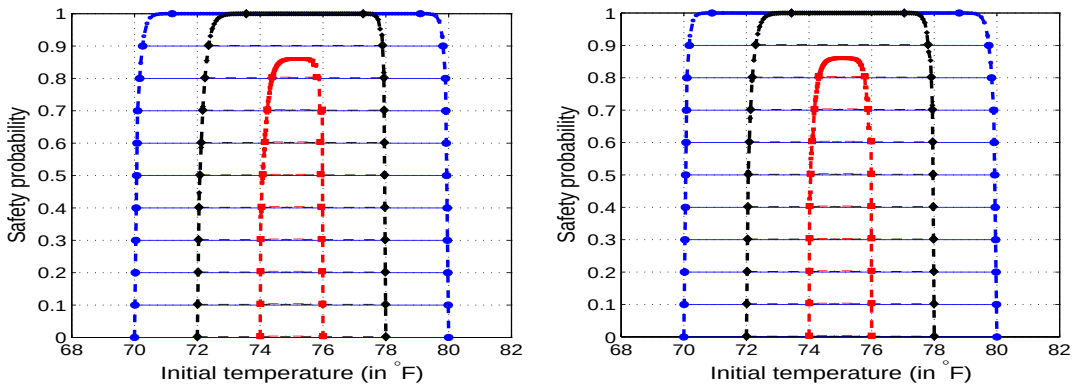

Figure 2.8: Maximal probabilistic safe sets: heater initially off (left) and on (right). Blue, black, and red colors refer to cases when the safe sets are A_1, A_2, and A_3, respectively.

Figure 2.9, this time interval gets progressively smaller moving from A_1 to A_2 and A_3.

It is interesting to note the behavior of the maximally safe policy corresponding to the safe set A_1 at $k = 575$ and $k = 580$. For example, for $k = 580$, the maximally safe policy for the `OFF` mode fluctuates between actions 0 and 1 when the temperature is around $75°F$. This is because the corresponding values taken by the function to be optimized in (2.19) (or, equivalently, in (2.20)) are almost equal for the two control actions. This feature may be exploited when we add to the safety problem a performance specification, as discussed in section 2.2.7.

Let us stress that the obtained results refer to the case of switching probabilities $\alpha = \beta = 0.8$. Different choices of switching probabilities may yield qualitatively different maximally safe policies.

2.3.2.2 Single Room, Regulation and Practical Stabilization Problem

In this section let us apply the methodology for reachability computations developed earlier to solve a practical stabilization problem. In particular, we will leverage the results on the infinite horizon case developed in section 2.2.4.

As in section 2.3.2, the transition control space is $\mathcal{U} = \{0, 1\}$, with the understanding that "1" means that a switching command is issued to the heater, "0" that no switching command is issued. As before, the reset control space is $\Sigma = \varnothing$, since

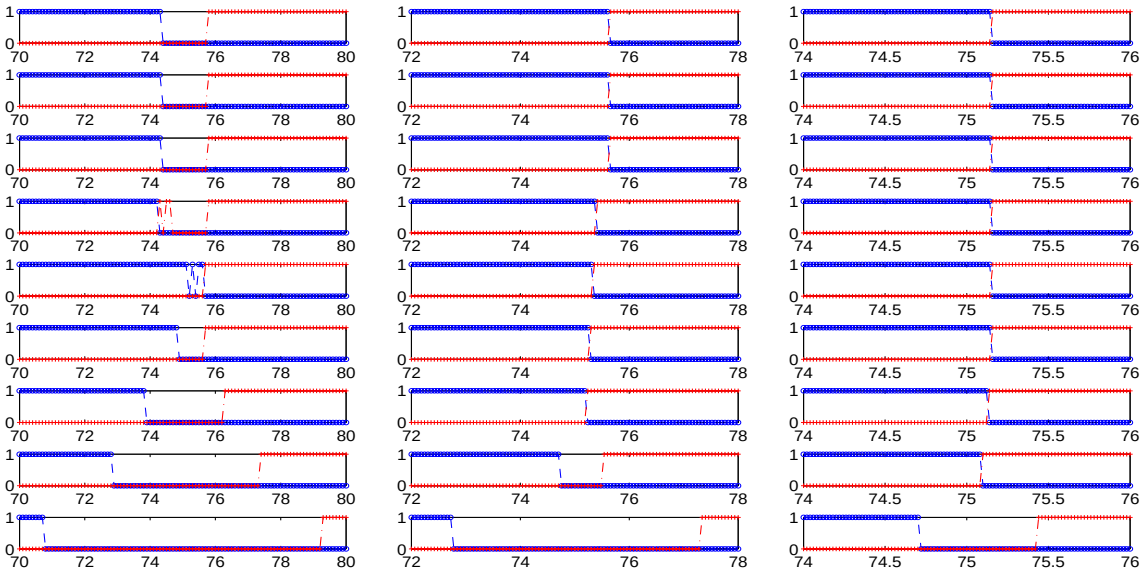

Figure 2.9: Maximally safe policy as a function of the temperature at times $k = 1$, 250, 500, 575, 580, 585, 590, 595, and 599 (from top to bottom) for the safe sets A_1 , A_2, and A_3 (from left to right). The darker (blue) circled line corresponds to the `OFF` mode and the lighter (red) crossed line corresponds to the `ON` mode.

the switching command to the heater is the only control input. The dynamics follow again the relations in section 2.3.2. The discrete transition kernel T_q is simplified as follows

$$T_q(q'|(q,x),0) = \begin{cases} 1, & q' = q \\ 0, & q' \neq q \end{cases}$$
$$T_q(q'|(q,x),1) = \begin{cases} \alpha, & q' \neq q \\ 1-\alpha, & q' = q \end{cases} \tag{2.44}$$

where $\alpha \in [0,1]$ represents the probability of switching from one mode to the other in one time-step.

The temperature is again measured in Fahrenheit degrees ($^\circ F$) and the time in minutes (min). The discretization time step Δ is chosen to be 1 min. The parameters are assigned the following values: $x_a = 10.5^\circ F$, $a/C = 0.1\ min^{-1}$, $r/C = 10^\circ F/min$, and $\nu = 0.33^\circ F$. The switching probability α in the above equation is set to be equal

to 0.8.

Let us consider the following regulation problem: determine a control law that maximizes the probability that the average room temperature x is driven close to $75°F$ in t *min* starting from any value in the set $(70, 80)°F$, with an admissible excursion of $\pm 1°F$ around $75°F$, and maintained within $75°F \pm 1°F$ thereafter. The variable t is the allowed time to steer the temperature to the desired region and can be specified by the user or can be chosen by the control designer. In the results reported next, let us consider any of the three $t = 150, 300, 450$. The implementation is done in MATLAB with a discretization step for the temperature coordinate equal to $0.05°F$.

The regulation problem can be reformulated as that of computing a maximally safe policy for a time varying *safe* set $A(k) = \mathcal{Q} \times \mathcal{X}(k)$, where $\mathcal{X}(k)$ shrinks from the region $(70, 80)°F$ towards the desired region $(74, 76)°F$ during the time interval $[0, t)$ *min*, and then keeps equal to $(74, 76)°F$ from time t *min* on.

The results discussed below refer to the following three different evolutions in time of the safe set for the temperature during the time interval $[0, t)$:

$$\mathcal{X}_1(k;t) = (70, 80),\ k \in [0, t)$$
$$\mathcal{X}_2(k;t) = \begin{cases} (70, 80),\ k \in [0, t/2) \\ (66 + \frac{8k}{t}, 84 - \frac{8k}{t}),\ k \in [t/2, t) \end{cases}$$
$$\mathcal{X}_3(k;t) = \left(70 + \frac{4k}{t}, 80 - \frac{4k}{t}\right),\ k \in [0, t).$$

Correspondingly, $A_i(\cdot;t) = \mathcal{Q} \times \mathcal{X}_i(\cdot;t)$, $i = 1, 2, 3$.

We determined the maximally safe Markov policies $\mu^\star(i, t)$, $i = 1, 2, 3$, for three different values of t: $t = 150$, $t = 300$, and $t = 450$. The maximally safe policies for $t = 150$, $t = 300$, and $t = 450$ are plotted, respectively, in Figures 2.10, 2.11, and 2.12. In each one of these figures, the plots in the first row refer to the safe set A_1, those in the second row to A_2, and those in the third row to A_3. The plots in the same column correspond to the `OFF` mode (plot on the left) and the `ON` mode (plot on the right). Each plot represents the value taken by the binary input u during the time horizon from 0 to 600 *min* (on the horizontal axis) as a function of the temperature (on the vertical axis). For any time instant $k \in [0, 600]$ only the corresponding safe

temperature range is considered. The value 0 ("do not switch") for u is plotted in gray, whereas the value 1 ("switch") is plotted in black. The maximally safe policies are expected to be time-varying during the time interval $[0, t)$. When $t = 450$ (Figure 2.12), the time variability is confined to the time interval $[378, 450]$ for A_1, whereas for A_2 and A_3 it extends to the intervals $[208, 450]$ and $[0, 450]$, respectively. This is due to the fact that in the last two cases, the safe set is time-varying over $[225, 450]$ and $[0, 450]$, respectively.

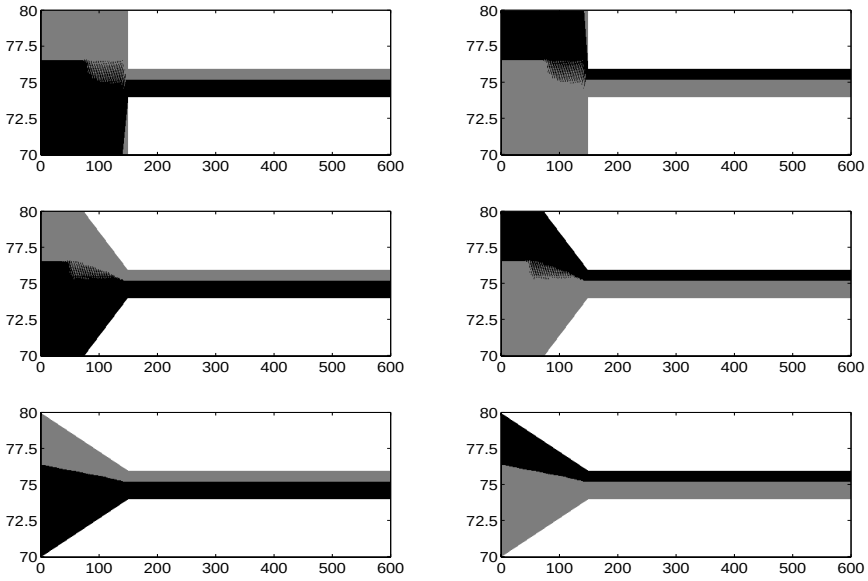

Figure 2.10: Maximally safe policy as a function of the temperature and time for the safe sets A_1 , A_2, and A_3 (from top to bottom). The left (right) column corresponds to the `OFF` (`ON`) mode. The darker (lighter) shade indicates that "switch" ("do not switch") is the recommended action ($t = 150$).

We also determined the probabilities $\mathrm{p}_{\pi}^{\mu^\star(i,t)}(A_i(\cdot;t))$, of remaining in the time-varying safe sets $A_i(\cdot;t)$, $i = 1, 2, 3$, $t = 150, 300, 450$, when the initial distribution π is uniform over $\mathcal{Q} \times (70, 80)$.

Not surprisingly, $\mathrm{p}_{\pi}^{\mu^\star(i,t)}(A_i(\cdot;t))$ increases with the transient length t, for any i. For each t, it instead remains almost the same for the three safe sets $A_i(\cdot;t)$, $i = 1, 2, 3$, with the value for $A_1(\cdot;t)$ only marginally higher than the others. For instance, when $t = 450$, the value is 0.831 for $A_1(\cdot;t)$, 0.822 for $A_2(\cdot;t)$, and 0.815 for $A_3(\cdot;t)$. This is easily explained by noting that, for each t, $A_3(k;t) \subseteq A_2(k;t) \subseteq A_1(k;t)$, $\forall k \in [0, \infty)$,

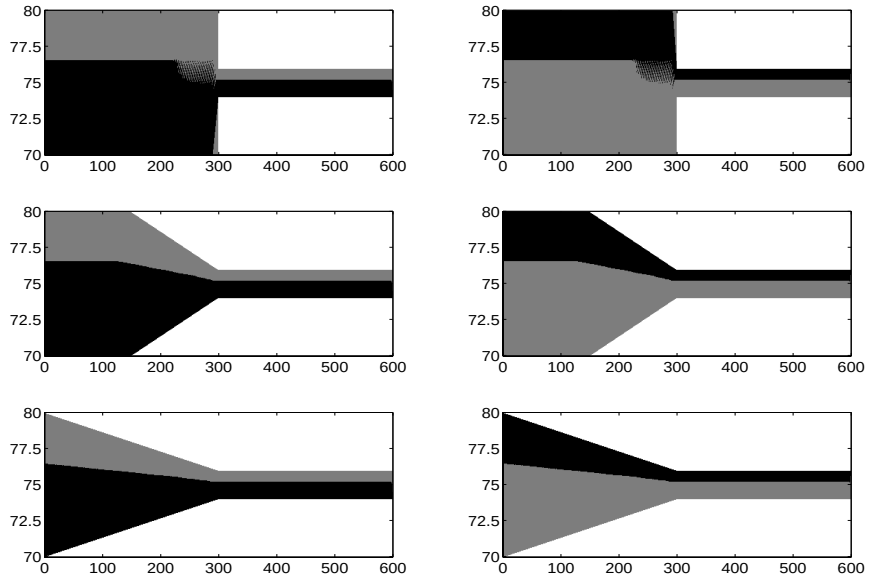

Figure 2.11: Maximally safe policy as a function of the temperature and time for the safe sets A_1 , A_2, and A_3 (from top to bottom). The left (right) column corresponds to the OFF (ON) mode. The darker (lighter) shade indicates that "switch" ("do not switch") is the recommended action ($t = 300$).

hence, $\mathrm{p}^{\mu}_{\pi}(A_3(\cdot;t)) \leq \mathrm{p}^{\mu}_{\pi}(A_2(\cdot;t)) \leq \mathrm{p}^{\mu}_{\pi}(A_1(\cdot;t))$, for any μ and π.

We determined also the maximal probabilistic safe sets $S^{\star}(\epsilon;i,t) = \bigcup_{q\in\{\texttt{OFF},\texttt{ON}\}}\{x \in \mathbb{R} : \mathrm{p}^{\mu^{\star}(i,t)}_{(q,x)}(A_i(\cdot;t)) \geq 1-\epsilon\}$ corresponding to different safety levels $1-\epsilon$. For any i, $S^{\star}(\epsilon;i,t)$ gets smaller as the required safety level $1-\epsilon$ grows, and get larger for higher values of t. For each t and ϵ, the maximal probabilistic safe sets are almost the same for the three safe sets. In Figure 2.13 we plotted the subset of $S^{\star}(\epsilon;i,t)$ corresponding to $q = \texttt{OFF}$ for $i = 1$ and $t = 150, 300, 450$ (the plots when $q = \texttt{ON}$ are similar).

2.3.2.3 Multiple Rooms Case

The temperature regulation problem described briefly in section 2.3.1 was extended to a multi-room scenario.

The objective is to maintain the temperature of r rooms within a certain range over some finite time horizon by heating one room at a time (single heater case). To this purpose let us devise a Markov policy that decides at each time instant which

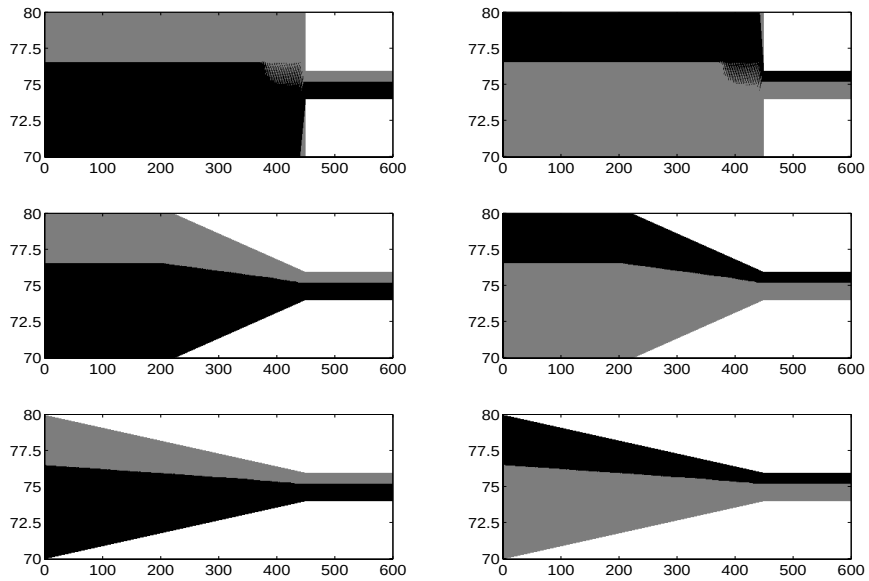

Figure 2.12: Maximally safe policy as a function of the temperature and time for the safe sets A_1 , A_2, and A_3 (from top to bottom). The left (right) column corresponds to the OFF (ON) mode. The darker (lighter) shade indicates that "switch" ("do not switch") is the recommended action ($t = 450$).

room should be heated based on the current value of the temperature in the r rooms. As before, this control design problem can be reformulated as a safety problem for the DTSHS model introduced in section 2.3.1, with the continuous state given by the r room temperatures and the discrete state given by the room that is being heated. The *safe* set is represented here by the desired temperature range for any discrete state.

Let us present the results obtained in the $r = 2$ rooms case. In this set of simulations, the temperature is measured in degrees Celsius and one discrete time unit corresponds to $\Delta = 10$ minutes. The discrete time horizon is $[0, N]$ with $N = 60$, which corresponds to an effective length of $N\Delta = 600$ minutes.

The discrete state space is $\mathcal{Q} = \{\mathtt{ON}_1, \mathtt{ON}_2, \mathtt{OFF}\}$ and the continuous state space is $\mathbb{R}^2$ in each mode $q \in \mathcal{Q}$ (in other words, the coordinates are the temperature level in each of the two rooms). The desired temperature interval is [17.5, 22] in both rooms. Thus, the safe set A is given by $A = \mathcal{Q} \times A_x$ with $A_x := [17.5, 22] \times [17.5, 22]$. The parameters values in equation (2.38) are set equal to: $x_a = 6$, $b_1 = b_2 = 0.0167$,

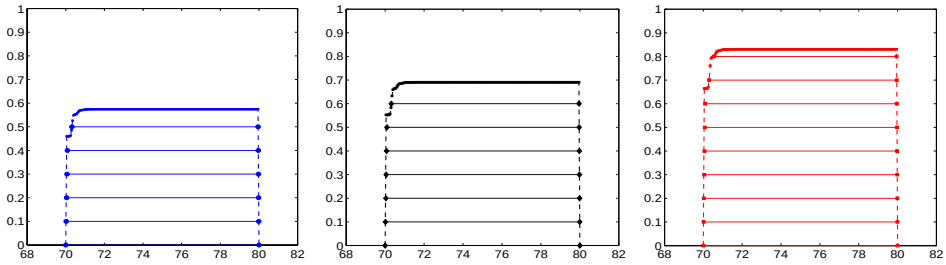

Figure 2.13: Maximal probabilistic safe sets for $\mathcal{X}_1(\cdot;t)$ when $t = 150, 300, 450$ (from left to right) and the heater is initially off. On the vertical axis is reported the safety level, on the horizontal axis the temperature values.

$a_{12} = a_{21} = 0.022$, $c_1 = 0.8$, $c_2 = 0.9333$, and $\nu^2 = 0.06$.[3] In this implementation, the temperature interval [17.5, 22] was discretized with a grid size 0.1, resulting in a uniform gridding of the set A_x into 45 intervals along each coordinate axis. Recall that in sections 2.2.8.1 and 2.2.8.2 it was shown that, under rather weak regularity assumptions on the stochastic kernels, as the continuous state gridding size goes to zero, the numerical solution to the dynamic programming equations tends to the actual solution with an explicitly computable convergence rate.

The transition control input takes on values in $\mathcal{U} = \{\mathtt{SW}_1, \mathtt{SW}_2, \mathtt{SW}_{\mathtt{OFF}}\}$ and affects the evolution of the controlled Markov chain governing the discrete transitions of the DTSHS model (see Figure 2.5). Let us suppose that when a command of commuting from a node to a different one is issued, then the prescribed transition actually occurs with a probability 0.8, whereas with probability 0.1 the situation remains unchanged (which models a delay) and with probability 0.1 a transition to the third, non-recommended node, occurs (which models a faulty behavior). Instead, when a command of remaining in the current node is issued, this happens with probability 1. These specifications can be precisely formalized by appropriately defining the controlled Markov chain transition probabilities $\{\alpha_{qq'}(u),\ q, q' \in \mathcal{Q}\}$ introduced in (2.40), for any $u \in \mathcal{U}$. For instance, for $u = \mathtt{SW}_1$, $\alpha_{\mathtt{ON}_1\mathtt{ON}_1}(\mathtt{SW}_1) = 1$, $\alpha_{\mathtt{ON}_2\mathtt{ON}_1}(\mathtt{SW}_1) = 0.8$, $\alpha_{\mathtt{ON}_2\mathtt{ON}_2}(\mathtt{SW}_1) = 0.1$, $\alpha_{\mathtt{OFFON}_1}(\mathtt{SW}_1) = 0.8$, and $\alpha_{\mathtt{OFFOFF}}(\mathtt{SW}_1) = 0.1$, the other probabilities $\alpha_{qq'}(\mathtt{SW}_1)$ being determined by the normalization condition $\sum_{q' \in \mathcal{Q}} \alpha_{qq'}(\mathtt{SW}_1) = 1$,

[3] The smaller dimension of the variance of the noise is also due to the new temperature range, which is scaled down as it is expressed in Celsius degrees, rather than according to the Fahrenheit scale.

$q \in \mathcal{Q}$.

The dynamic programming recursion described in section 2.2.3 can be used to compute a maximally safe policy $\mu^\star = (\mu_0^\star, \mu_1^\star, \ldots, \mu_{N-1}^\star)$, $\mu_k^\star : \mathcal{S} \to \mathcal{U}$, $k = 0, 1, \ldots, N-1$, and the maximal probabilistic safe sets $S^\star(\epsilon)$, $\epsilon \in [0, 1]$. In the implementation, the multiplicative setup was chosen, and computations were performed in MATLAB.
Figure 2.14 shows some 'optimal' sample paths of the continuous state component of the DTSHS executions associated with the maximally safe policy $\mu^\star$ and different initial conditions. The initial conditions were chosen at random, according to the uniform distribution, over the safe set A. Note that, as expected, most of the temperature sample paths tend toward the central area of A_x, and those sample paths that exit A_x correspond to initial conditions close to its boundary. This is due partly to the delay in the commutations, and partly to the noise affecting the temperature dynamics.
In Figure 2.15,[4] let us represent the component of the maximal probabilistic safe set $S^\star(\epsilon)$ associated with the discrete state $\texttt{OFF}$, that is $\{x \in \mathbb{R}^2 : (\texttt{OFF}, x) \in S^\star(\epsilon)\} \subseteq A_x$, for different safety levels $1 - \epsilon$. The plots corresponding to the discrete modes $\texttt{ON}_1$ and $\texttt{ON}_2$ are similar. As expected, the maximal probabilistic safe sets get smaller as the required safety level $1 - \epsilon$ grows. Also, their shape reveals some asymmetry due to the structure of the temperature dynamics. Because of the low value of the ambient temperature ($x_a = 6$), the temperature tends naturally to decrease (see equation (2.38)).

The values taken by function $\mu_0^\star : \mathcal{S} \to \mathcal{U}$ over the set A_x when $q = \texttt{OFF}$ are plotted in Figure 2.16.[5] $\mu_0^\star(\texttt{OFF}, x)$ is the maximally safe transition control input issued at time $k = 0$ when $\mathbf{s}(0) = (\texttt{OFF}, x)$. The maximally safe controls for the other time steps k within the horizon $[0, N]$ are indeed really much like the one in Figure 2.16, except for the very final time steps. This shows that a *stationary control policy* would be nearly maximally safe in this case. Different choices of the discrete transition probabilities (2.40) may yield qualitatively different maximally safe policies.

[4] Notice that this figure coincides with the bottom right one in 2.18, modulo properly modified parameters of the dynamics.

[5] Notice that this figure coincides with the bottom right one in 2.19, modulo properly modified parameters of the dynamics.

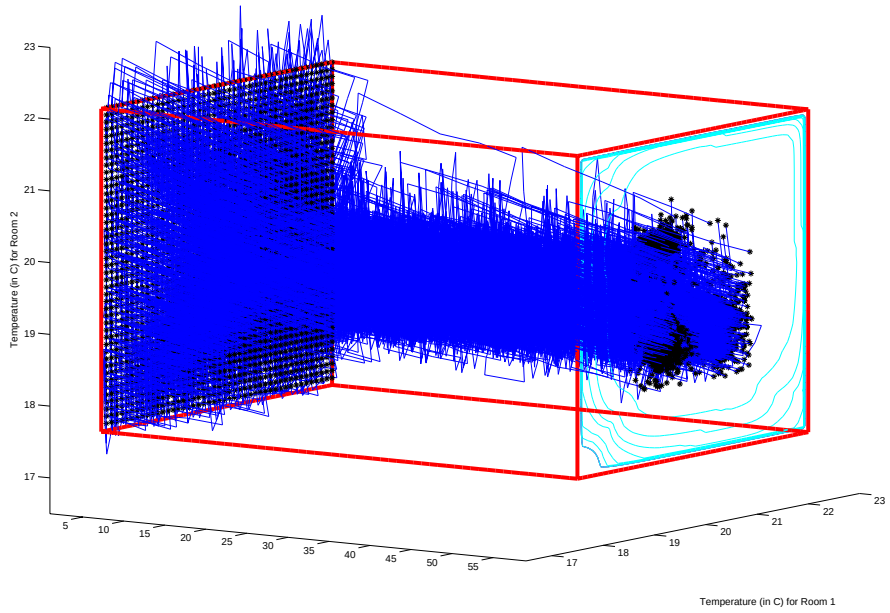

Figure 2.14: Sample paths of the two rooms temperatures for executions corresponding to different initial conditions, under the same maximally safe policy.

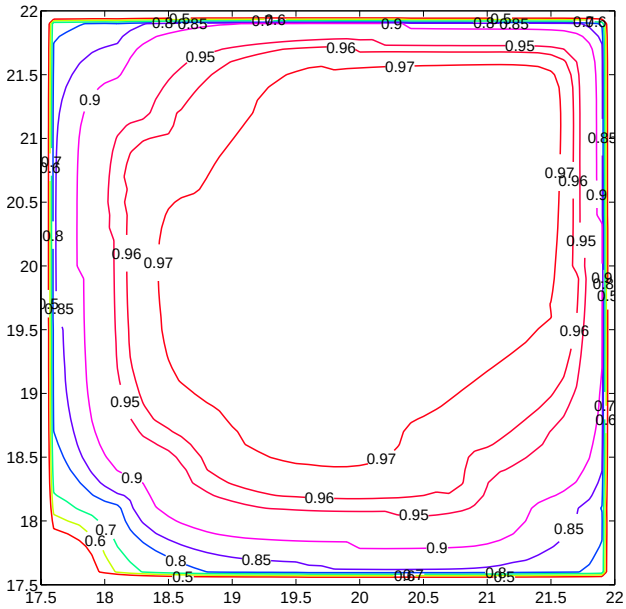

Figure 2.15: Maximal probabilistic safe sets corresponding to different safety levels (0.5, 0.6, 0.7, 0.8, 0.85, 0.9, 0.95,0.96, and 0.97) within the discrete state `OFF`. The temperature of room 1 is reported on the horizontal axis, and that of room 2 on the vertical axis.

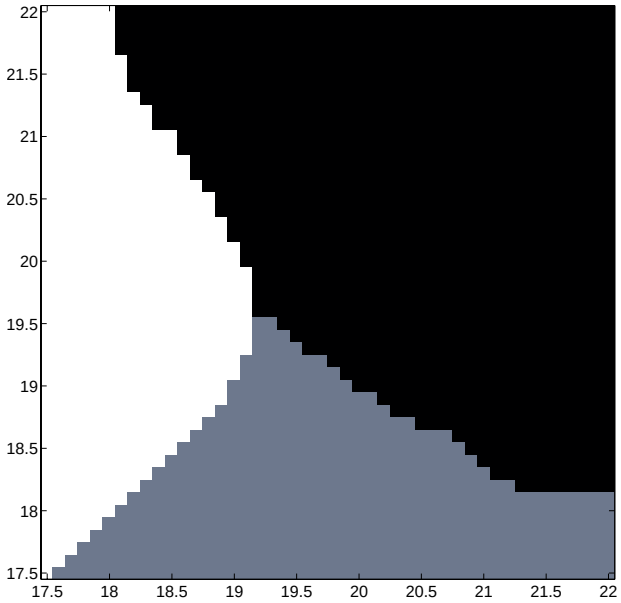

Figure 2.16: Value taken by the maximally safe transition input at time $k = 0$ over the set of desired temperatures, when the heating system is in the OFF mode. The temperature of room 1 is reported on the horizontal axis, and that of room 2 on the vertical axis. The colors *black*, *white*, and *grey* respectively stand for the transition input command $\mathtt{SW}_{\mathtt{OFF}}$, $\mathtt{SW}_1$, and $\mathtt{SW}_2$.

2.3.3 Numerical Approximations

In this section let us present the results of a computational study for the multi-room heating benchmark where the convergence of quantities computed by the discretization scheme proposed in section 2.2.8.1 and analyzed in section 2.2.8.2 are numerically studied.

The benchmark in [Fehnker and Ivančić, 2004] deals with the problem of regulating temperature in a house with n rooms via m heaters. In this section let us focus on $m = 1$, the single heater case, formally introduced in section 2.3.1, and which notations we shall mostly adhere to in the rest of the work.

Aiming at highlighting the structure of the controller more explicitly than in section 2.3.1, let us introduce the following formalism. The heater is controlled by a thermostat that is prone to delay and failures in switching the heater between one room to another and between the "on" and "off" status: the effect of these control actions on the discrete state transitions is specified by a finite-state, finite-action, controlled Markov chain which is independent of the continuous state, that is, $T_q : \mathcal{Q} \times \mathcal{Q} \times \mathcal{A} \to [0, 1]$. One can easily check that the number of possible discrete states is $n + 1$ and the maximum number of available control actions is $n(n + 1) + 1$ (notice that this number is smaller than that in (2.40), because we cluster the "do not switch" action in each mode into a single action). Let us define the compact safe set to be

$$A = \bigcup_{q \in \{1,\ldots,(n+1)\}} \bigcup_{i \in \{1,\ldots,n\}} \{q\} \times \{i\} \times [x^q_{li}, x^q_{ui}],$$

where x^q_{ui} and x^q_{ui} specify the lower and upper limits for the desired temperature in room i for discrete state q. For simplicity, these are assumed to be independent of i and q. Let us now describe the discretization procedure as follows: let us adopt a uniform partitioning of the set $[x^q_{li}, x^q_{ui}]$ into m disjoint intervals each of size $\varkappa = (x^q_{ui} - x^q_{li})/m$. Therefore, $[x^q_{li}, x^q_{ui}] = [x^q_{li}, x^q_{li} + \varkappa) \bigcup, \ldots, \bigcup, [x^q_{li} + (m-1)\varkappa, x^q_{ui}]$. The value of the temperature in room i for the discrete state q is defined by $e^q_{r_i i} = x^q_{li} + (r_i - 1)\varkappa$, where $r_i \in \{1, \ldots, m\}$. Define $r = [r_1, \ldots, r_n]^T$. Let us pick $v^q_r = (q, [e^q_{r_1 1}, \ldots, e^q_{r_n n}]^T)$ as hybrid state value. Thus, the set of all discrete values for the hybrid state is $\mathcal{G} = \{v^q_r, r = [r_1, \ldots, r_n]^T; r_i = 1, \ldots, m; i = 1, \ldots, n; q = 1, \ldots, (n + 1)\}$. Let $\mathcal{N}(\cdot; \eta, \sigma^2)$

denote the probability measure over $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$ associated with a Gaussian density function with mean η and variance σ^2. Then the stochastic kernel $T_s(ds'|v_r^q, a)$ that is used in the discretized dynamic programming equations (2.30) be defined as follows:

$$T_s(ds'|v_r^q, a) = T_x(dx'|v_r^q, a)T_q(q'|q, a),$$

for $v_r^q \in \mathcal{G}$, $a \in \mathcal{A}$, and $s' \in \mathcal{S}$. Here, $T_x(\cdot|v_r^q, a) = \mathcal{N}(\cdot; \mu_r^q, \nu^2 I_n)$, I_n being the identity matrix of size n, $\mu_r^q = [\mu_{r_1}^q, \ldots, \mu_{r_n}^q]^T$ and $\mu_{r_i}^q = e_{r_i}^q + \Big(b_i(x_a - e_{r_i}^q) + \sum_{i \neq j} a_{i,j}(e_{r_j}^q - e_{r_i}^q) + c_i h_i\Big)\Delta$. It is easy to check that $T_x(dx'|v_r^q, a)$ and $T_q(q'|q, a)$ satisfy the Assumption 8.

Convergence Properties. Let us first analyze the convergence properties of the discretization scheme for the case when $n = 2$ (two rooms). The number of modes is 3 and maximum number of allowable control actions is 7, as shown in Figure 2.17(a). The computations are performed for the safe set $A = \bigcup_{q \in \{1,2,3\}} \bigcup_{i \in \{1,2\}} \{q\} \times \{i\} \times [17.5, 22]^o$C. The size of time interval is $\Delta = 1/15$ and the number of intervals is $N = 60$. The parameters values in equation (2.38) are: $x_a = 6$, $b_1 = b_2 = 0.25, a_{12} = a_{21} = 0.33, c_1 = 12, c_2 = 14$ and $\nu^2 = 0.9$.[6] All the parameters should be interpreted in appropriate units. For each control action by the thermostat that elicits a transition between two different modes of the heater, the transition happens with probability 0.8. The remaining 0.2 probability is divided evenly between the "do nothing" transition that models the delay and the transition to the third, non-recommended mode that models a faulty behavior.

The computations of the solutions $\widehat{V}_0^*$ to the approximated DP equations in (2.30) were performed for four discretization levels: $m \in \{9, 18, 36, 45\}$. Following section 2.2.1, let us define the *approximately maximal probabilistic safe sets* $\widehat{S}^*(\epsilon)$ with safety level $(1-\epsilon)$ as $\widehat{S}^*(\epsilon) = \{s \in \mathcal{S} : \widehat{V}_0^*(s) \geq (1-\epsilon)\}$. Figure 2.18 shows the approximately maximal safe sets when the initial discrete state is `OFF`, and corresponding to different safety levels. As expected, the maximal safe sets get smaller as the required safety level increases. Furthermore, as the discretization level decreases, the maximal safe sets tend to graphically converge: this visually confirms the numerical convergence of

[6] In this set of simulations, lat us assume a time-discretization interval that is smaller than the one assumed in the previous case studies. This is because of the improved needed precision. Hence, the parameters are properly rescaled, for the sake of uniformity with the other computations.

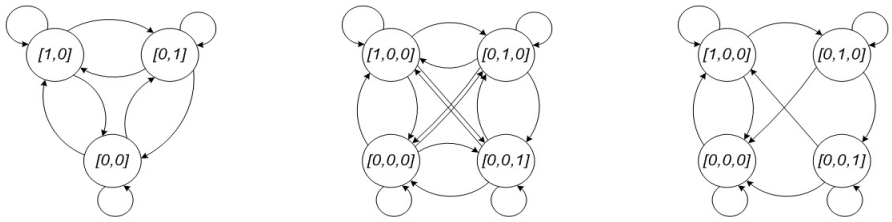

Figure 2.17: (a) Maximum available control actions for $n = 2$. (b) Maximum available control actions for $n = 3$. (c) Reduced number of available control actions for $n = 3$. The discrete states are assigned numbers clockwise starting from the top-left state.

the proposed discretization scheme.

The optimal control actions for the case when the initial discrete state is `OFF` are plotted in Figure 2.19 for the four discretization levels and $k = 1$. The optimal actions at finer resolution were obtained from that of coarser resolution by nearest neighbor interpolation. It can be noticed that the regions of optimal recommended actions become more well-formed and again visually converge as the discretization step decreases.

2.3.4 Mitigating the Curse of Dimensionality

To conclude this part dedicated to the computational aspects of the proposed methodologies, let us propose possible improvements in the implementation. It is of interest to understand when the underlying structure of the DTSHS can be exploited to implement the DP algorithm in a computationally efficient manner. More precisely, let us reinterpret some ideas proposed in [Bertsekas, 1982] and [Kveton *et al.*, 2006] within the hybrid systems framework to suggest that particular structural properties of the problem, such as its decentralized nature, may be exploited to obtain a more compact state representation and efficient implementation of the computations involved in the solution of the DP. This feature may at least in part mitigate the curse of dimensionality the DP is doomed to, but which similarly affects the other related approaches to the problem in the literature [Mitchell and Templeton, 2005].
For a more thorough discussion and further references, the reader is referred to section 2.5 (Computational Improvements).

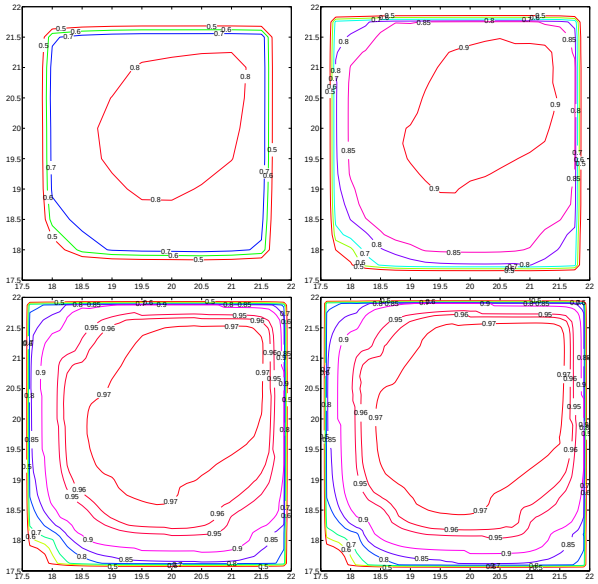

Figure 2.18: Maximal probabilistic safe sets corresponding to safety levels: $0.5, 0.6, 0.7, 0.8, 0.85, 0.9, 0.95, 0.96$ and 0.97 for the case $n = 2$ and initial discrete state OFF. In going from left-to-right and top-to-bottom, the plots correspond to discretization levels of $9, 18, 36$ and 45 respectively.

2.3.4.1 Scaling to Higher Dimensions

Let us present the results from the three-room, one heater benchmark case. For this case, the number of continuous states is $n = 3$, the number of discrete states is 4 and maximum number of allowable actions is 13, as shown in Figure 2.17(b). The safe set is specified to be $A = \bigcup_{q\in\{1,2,3,4\}} \bigcup_{i\in\{1,2,3\}} \{q\} \times \{i\} \times [17.5, 22]^o\text{C}$. The size of time interval is $\Delta = 1/15$ and the number of intervals is $N = 60$. The parameters values in equation (2.38) are: $a_{12} = a_{21} = 0.80$, $a_{13} = a_{31} = 0.60$, $a_{23} = a_{32} = 0.70$, $x_a = 6$, $b = [0.30, 0.20, 0.30]^T$, $c = [12.00, 14.00, 12.00]^T$ and $\nu^2 = 0.33$. Similar to the two-room case, the effect of control actions is described by a controlled Markov chain. The computation of the DP algorithm was performed for the discretization level $m = 18$. Figure 2.20 shows the maximal safe sets corresponding to the safety level $1 - \epsilon = 0.95$, at different times. As expected, as the number of steps-to-go increases, the size of the safe sets also decreases. It is of interest to compare the effect of number of available control actions on the size of the maximal safe set. In order to study this, we performed the DP computations for the three-room, one heater example

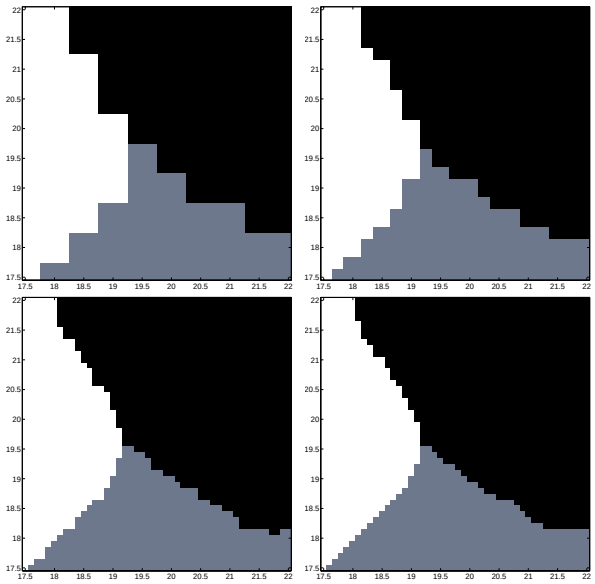

Figure 2.19: Maximally safe actions for the case $n = 2$, initial discrete state `OFF` and $k = 1$. In going from left-to-right and top-to-bottom, the plots correspond to discretization levels of $9, 18, 36$ and 45 respectively. The colors *black*, *white* and *grey* respectively stand for "do nothing," "switch heater to room 1" and "switch heater to room 2" actions.

for the reduced set of actions shown in Figure 2.17(c). The resulting maximal safe sets corresponding to safety level $1-\epsilon = 0.95$ are shown in Figure 2.21. Let us observe that the maximal safe set is becomes very small and eventually decreases to the empty set as the number of steps-to-go increases. Let us finally notice an important structural property of the benchmark, namely the conditional independence of the continuous stochastic kernel: $T_x(\bar{x}|v_r^q, a) = T_x(\bar{x}_1|v_r^q, a) \times \ldots \times T_x(\bar{x}_n|v_r^q, a)$. This enables us to efficiently compute the state transition probabilities, as well as to propagate the calculations of the DP scheme in a "distributed" fashion. Please refer to section 2.5 for further details.

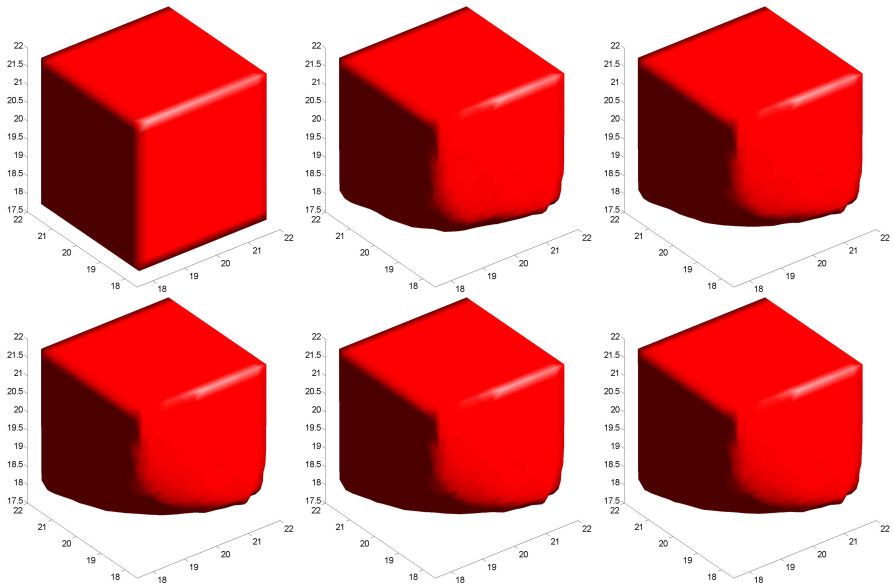

Figure 2.20: Maximal probabilistic safe sets corresponding to a safety level of 0.95 for the case when $n = 3$ and initial discrete state is `OFF`. Available control actions are shown in Figure 2.17(b). In going from left to right and top to bottom, the plots correspond to $k = 60, 55, 50, 40, 20$ and 1.

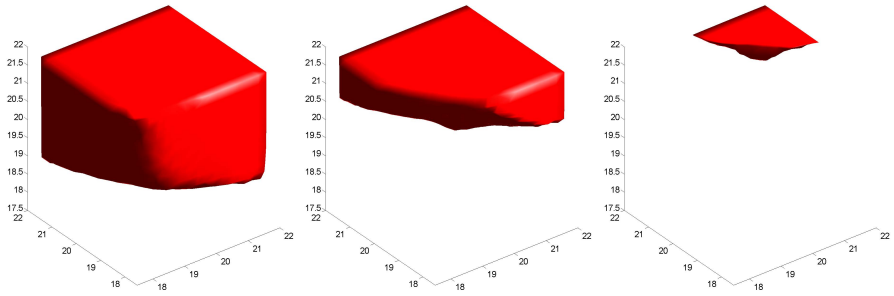

Figure 2.21: Maximal probabilistic safe sets corresponding to the safety level 0.95 for the case $n = 3$ and initial discrete state `OFF`. The reduced set of available control actions is shown in Figure 2.17(c). In going from left-to-right, the plots correspond to $k = 55, 50$ and 45. The safe set for $k = 60$ is same as the corresponding safe set in Figure 2.20.

2.4 Applications

In this section, an application of the concept that has been introduced in this chapter will be developed as a case study taken from Biology.

Briefly pausing from this task, let us remark that, more generally, a wealth of possible applications can be envisioned for a general concept such as that of reachability, as well as that of safety. The *regulation problem* introduced in section 2.2.5 further connects with important possible applications related to the objective of system stabilization. In section 2.2.6 a potential direct application of the concept in *fault tolerant* control problems has been investigated. Furthermore, provided a game-theoretical setup is developed from the current framework (which is a topic of future research), a number of other engineering problems could also be modeled and studied. The computational study presented in section 2.3.4 hinted at the analysis of distributed control strategies in networks as a possible target for the proposed methodology. The author has looked, for the sake of instance, at the use of reachability and safety analysis for *buildings automation* applications. Other areas that appear to be easy and fruitful target for the techniques presented in this thesis are *disease spread analysis and control*, or power management studies in *communication networks*.

In the following, let us instead concentrate on a methodological study in System Biology. Starting from a general concept, that of optimality, let us propose a way to interpret this notion in a study of survival analysis. This approach is further framed within a particular biological case study, by reinterpreting survival as a safety specification, and optimality is obtained and interpreted by solving a related optimal control problem. This solution is associated to specific production mechanisms and stress-response behaviors observed in nature.

Mathematical Biology and Systems Biology. The application of the qualitative theory of ordinary differential equations in the biological sciences has a long and successful history [Rubinow, 1975]. Fibonacci, Lotka, Vehrlust, Volterra and Maynard Smith are just a few examples of mathematicians that have left a groundwork contribution in biological modeling and analysis, furthering the understanding of such real-world systems. From a taxonomic perspective, the application of mathematical models in Biology can be done at different levels: from the large scale (population dynamics, ecology, epidemiology), through the human level (physiology, rehabilitation and medicine), to the small scale (cellular, molecular and genetic).

More recently, the attention has been focused especially at the cellular and molecular level, with the study of genetic pathways, metabolic networks, protein-protein interactions, signalling pathways. This research has been made possible by the availability of data in a thus far unexperienced quantity. Over-employed is "the post-genomic era" slogan, which is often used to describe the new frontiers opened by the available knowledge of quantitative information of entities–and their functions–at the genomic level. This wealth of data comes, in particular, in the form of microarrays, databases, imaging, vision and high-throughput experiments and by the use of reporter systems.

The field of Systems Biology leverages the availability of data, and responds to the need for proper mathematical models to analyze, interpret and understand it. Furthermore, it has a clear focus on the *system* properties of a biological entity [Kitano, 2002]: the motives and modules that it builds [Rao and Arkin, 2001; Wolf and Arkin, 2003; Voigt *et al.*, 2005], their interconnections, and the network structures that emerge from these clusters [Strogatz, 2001].

From a different categorical perspective, if we look at the complete spectrum of possible mathematical and formal models in biology, then it is possible to single out, at the two ends, either purely discrete models, or purely continuous ones.
Examples of discrete models are, for the sake of instance, finite-state automata, structures from graph theory, Boolean and Petri nets. The advantage to exploit this discrete modeling framework is that a lot of information can be extracted directly from its structure (connectivity, cycles, dependencies, functions), the algorithms that can be used on it and their associated complexity.

Continuous models, at the other end of the spectrum, are capable of modeling the evolution of concentrations (which, for example, can be derived from stoichiometric reactions), the rates of these dynamics, and the possible dependence on spatial components (distributed-parameters systems). The advantages related to continuous models are their precision, the possibility to encompass multiple time scales, the possible study of their robustness properties, and their sensitivity analysis. More generally, they benefit from a suite of techniques, offered by Systems and Control Theory, that helps in investigating their properties.

The observation suggests that many biological instances are naturally made up of systems with interacting continuous and discrete components. One notable example is the *cell cycle* [Novak and Tyson, 1997], where phases of continuous growth are interleaved with discrete, logical conditions that influence one out of a number of possible outcomes (continued growth, apoptosis, or quiescence for instance). From a different level, biological systems often present heterogeneity coming from the interplay of behaviors happening at different levels (for instance, cellular vs. environmental).

It is intuitively desirable then to introduce a formal approach for integrating the previous two modeling efforts into a systematic whole. Hybrid Systems theory offers this opportunity. It is then not surprising that the use of Hybrid System models has been recently advocated in the Systems Biology literature [Rao and Arkin, 2001; De Jong, 2002; Lincoln and Tiwari, 2004; Cho *et al.*, 2005].

An undeniable feature of biological models, especially at the small-scale level, is the presence of noise, and the lack of complete and precise knowledge of the system. This uncertainty of the actual dynamics often comes from the quality of the data which, given its origin, is measured with errors. Furthermore, often the data is available in a coarse form, i.e. it is difficult to single out the signals or features of interest among the noisy and information-crammed observations. The role of noise actually goes beyond this simple and intuitive negative connotation. It is argued that it often plays a constructive role in biological phenomena [McAdams and Arkin, 1997; McAdams and Arkin, 1998; Hasty *et al.*, 1999; McAdams and Arkin, 1999; Rao *et al.*, 2002; Blake *et al.*, 2003]. Also, the nature of certain phenomena, both continuous and discrete, appears to be ruled by chance, especially when a limited number of entities are interacting [Gillespie, 1976; Gillespie, 1977].

The above arguments lead to the suggestion that the use of *Stochastic* Hybrid Systems is a natural choice of modeling and analysis framework for a wide class of biological systems.

2.4.1 Survival Analysis of Bacillus subtilis

The investigation of the stress response network of *Bacillus subtilis* ATCC 6633 offers a detailed explanation of how the bacterium reacts to competitive environmental conditions, among the many options, by producing the antibiotic subtilin in order to directly suppress other cells while getting immunized [Msadek, 1999]. The antibiotic production is just one possible action, among a network of potential responses that the bacterium can employ to survive. The mechanisms of the generation of the antibiotic subtilin are fairly well understood and described by a genetic and protein pathway that involves some non-deterministic interplay between its quantities. The interplay is non-deterministic because of the noise in the environment and the partial knowledge we have of the whole production mechanism. More precisely, the presence of switching modes exhibits the activation/deactivation of certain genes and the increase/decrease in production of the corresponding proteins. Furthermore, the dynamics and the triggering conditions are intrinsically not perfectly known, and hence prone to be modeled in probabilistic terms. Furthermore, if we interpret any response mechanism as a "control action," then it may well be that a randomized sort of control could be employed by the bacterium [Wolf *et al.*, 2005].

The concept of *optimum* is common and shared between engineering and biological systems. Prior to its employment, though, it is necessary to answer some fundamental questions: optimality with respect to what? And at what level? And under which constraints? The biology literature offers numerous examples where optimality appears to regulate a certain behavior, or to explain the properties of a particular entity. [Rosen, 1967] presents probably the first attempt, rather qualitative in nature but nevertheless very stimulating, to systematically frame the concept of optimality in biology. At various "levels," a number of observed biological phenomena are interpreted as a solution of certain optimal problems: the vascular system, at a physiological level, is deemed responsible of optimally transporting

the fluid within a body, and of optimally covering its interior with its branching network; the phenomenon of homeostasis hints at how robustness and tolerance in cellular and physiological entities is achieved by minimizing the error offset from certain references; similarly for a number of biological regulators, instance of optimal adaptive systems; the *allometric law*, dealing with form functionals, describes the evolution of shapes in as a proportion-preserving phenomenon, at the cellular and physiological level. More recently, [Segre *et al.*, 2002] employs an optimality interpretation in the context of metabolic networks, where the objective or to maximize the survival analysis at the cellular level. Contributions in metabolic flux analysis also exploit perspective [Klapa, 2007]. Furthermore, [Weibull, 1995; Wolf *et al.*, 2005] look at dynamical game theory as a means to think about optimality in the context of evolution, both at a population level, as well as at a cellular one. Along the same lines, theoretical questions about evolution and optimality are raised by [Maynard Smith, 1982]—in his work optimality is related to the notion of stability of a strategy.

On the other hand, many notable instances from the same biological domain caution that the abuse of this notion may yield to incorrect conclusions for such systems.

According to the general tenets of evolution, it can be cogently argued that *a biological structure is "optimal" (at least locally and temporarily) because it has survived evolution under the pressure of natural selection.* It is important to stress the local value of the statement, as well as its "non uniformity" in time. A few comments on the above statement are due. While such notion is generally accepted by the proponents of the so-called *adaptationist* approach to evolutionary sciences (notable advocates are Maynard Smith, Dawkins, Pinker and a number of ethologists), other scientists contend its over-accentuation of the power of natural selection [Gould and Lewontin, 1979] and its neglect of constraints (genetic, as well as ecological). In an attempt to reconcile these two takes, in [Duprè, 1987] the difference between *optimality as a fact*, rather than *as a heuristic*, is highlighted. While the first concept relates possible optimality traits in organisms from the role of natural selection in evolution, the second claims that if a trait is optimal, then it evolved by natural selection. Clearly the first position is less pervasive than the second. Furthermore, it does not hint at any pos-

sible teleological argument, which the present work has no intention to embrace. To quote [Duprè, 1987, p.2, italic added], optimality as a fact may emerge as "evolution ... is a process by which organisms have come to adopt the best available *strategy* for dealing with the exigencies of their environment." Notice that a strategy can be just partly reflected by phenotypic traits, upon which natural selection plays a role. Hence, by the factual tenet, a strategy may display some optimality properties.

With regards to the "level," that is the focus of the selection forces, we assume it to be the single species. Again, this is arguable, in that it may be postulated that instead the whole group, or the kin, and their corresponding fitness, should be the objective (the "unit") of study. To hunt away further possible criticism of this last point, and then of the validity and generality of the above claim, let us also stress that, along with that of survival, the objective of offspring yield (that is, its maximization and nurture) should be accounted for. However, we shall disregard this second objective. This is because of the nature of the present work, which focuses purely on the stress response network, rather than on other, more general behaviors related to survival.

In our study, the above reasoning translates into postulating that the functioning of the subtilin production pathway follows certain criteria and levels of optimality. In this context optimality is intended as a measure of personal fitness or, in the particular instance, of survival of the single bacterium. In particular, one would expect that the activation/deactivation of the production path in the network happens "optimally" in the above sense.

Bacillus subtilis has been the object of much experimental work, analysis and synthesis because of its availability, accessibility, and easy manipulation. In particular, the subtilin production pathway has been deeply investigated, both at a genetic and at a signaling level (further details can be found in section 2.4.2).

In this work, let us look at a recently developed dynamical model for the genetic network describing the biosynthesis of the lantibiotic subtilin [Hu *et al.*, 2004] and propose a few improvements and modifications to the model to bring it in line with newer evidence reported in the literature [Tjalsma *et al.*, 2004; Stein, 2005]. A system that presents partially decoupled high-level dynamics (those dealing with the population size and the nutrient level) and low-level ones (those describing the mechanism

of production of subtilin by a single cell) is obtained. The high-level model is nonlinear and deterministic, while the low-level one is hybrid, affine and probabilistic. The model, while limited in scope and properly reduced to a workable abstraction, has to be looked at and interpreted as a framework for the methodological procedure presented in this work. This endeavor does not claim to shed new light into the *dynamics* of the antibiotic pathway, but rather on its structure. More specifically, the effort is focused on the artificial synthesis of the control mechanism for the subtilin production, and on an interpretation of this mechanism in terms of a certain optimality criterion.

The system in its entirety can be interpreted (and therefore modeled) as a stochastic hybrid system. Upon this modeling setup, let us develop our study in two steps:

1. reinterpret survival analysis as a probabilistic safety problem;
2. leverage (stochastic) optimal control theory to address this problem in terms of optimality.

This approach allows one to study the survival of the single *B. subtilis* cell as a probabilistic, decentralized safety specification problem. It is "probabilistic" because of the inevitably stochastic dynamics. It is "decentralized" because each entity, while optimizing for its own fitness, does not communicate with all the competitors, nor has knowledge of their single actions, but only receives an averaged feedback from the environment (this signal could be due to the process of *quorum sensing* [Tjalsma *et al.*, 2004], for instance). The solution of this survival problem may not then be globally optimal. Using the techniques developed in this thesis, let us reinterpret the above probabilistic safety problem as a (stochastic) optimal control one, where the controls are functions of the state-space and encode the subtilin production strategies. We argue that, under a proper choice of the survival function (i.e., the safety region), the solution of this stochastic and decentralized optimal control problem yields the location and the structure of the switching behaviors under study. We draw some fair comparisons on the outcomes of the above procedure based on a set of putative and spurious survival functions.

This methodological approach may shed light on the strategies employed by the organisms to improve their chances of survival, conditioned on the possibility of cod-

ifying such fitness into a cost function. More generally, let us argue that, in certain instances, a function in a genetic network may be synthesized by an optimization problem, where the objective is to maximize the probability of survival of the individual. Prospectively, matching these outcomes with the data in the literature allows concluding that the corresponding switching mechanisms in the subtilin production network may function with a degree of optimality, according to certain survival criteria. Moreover, this approach suggests that it is according to this survival interpretation, rather than other ad hoc characterizations, that the thresholds in the system could be specified and determined.

2.4.2 A Model for Antibiotic Synthesis

Resorting to a schematization proposed in [Hu *et al.*, 2004], based on a wealth of recent research [Banerjee and Hansen, 1988; Entian and de Vos, 1996; Kiesau *et al.*, 1997; Msadek, 1999; Stein *et al.*, 2002; Kobayashi *et al.*, 2003; Stein, 2005], it is possible to abstract the biosynthesis network into a five-dimensional model (see figure 2.22, taken from [Hu *et al.*, 2004]). The model encompasses two "global" variables (population and nutrient level) and three "local" ones (the concentration of the sigma factor SigH and of the two proteins SpaRK and SpaS).

In order to prune away details that may be uninteresting at this level, the presence of the peptide SpaS is equated to represent the actual antibiotic subtilin. In actuality, a proportional relationship has been observed between these two quantities. For the sake of simplicity the scheme disregards some of the components in the otherwise complex subtilin biosynthesis pathway, as well as some behavior only tangentially of interest at this level. For instance, in Figure 2.22 one can notice the presence peptides SpaF, Spae, SpaG and SpaI, which play a role for the immunity response [Klein and Entian, 1994; Stein *et al.*, 2003; Stein *et al.*, 2005]. This immunization action prevents the cell from perishing by producing its own antibiotic. Let us disregard their role for now.

Furthermore, for the sake of simplicity (and, as it shall be seen in the following, of computation), given the symmetric and repetitive structure of the dynamics of SpaRK and SpaS in their dependence on, respectively, SigH and SpaRK [Hu *et al.*,

2004], in the current case study let us disregard the evolution of SpaRK and just hypothesize a direct influence of SigH on SpaS, without much change in the final behaviors. Compare, for reference, the trajectory plots in [Hu *et al.*, 2004, Figure 2].

Ideally, a model that keeps track of each of the dynamics of a variable set of species in a certain environment may yield extremely precise results. However, this approach is discarded because of its sheer impracticality and because of the difficulty in its dynamical analysis. Hence, it is preferred to introduce a model that decouples high-level from low-level dynamics (see figure 2.23). The higher level, which encompasses the first two global variables (population level and food level), is deterministic and based on average dynamics. The lower one, involving the last three local coordinates (protein and sigma factor concentration levels), describes cellular processes and is made up of stochastic and switching dynamics. The reader may notice that the model is endowed with a decentralized structure (see figure 2.23).

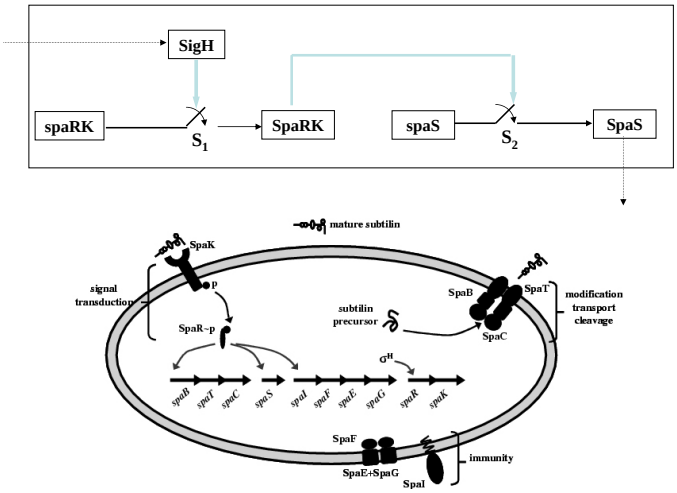

Figure 2.22: The subtilin biosynthesis pathway. The top figure represents a scheme where the relationship among the most important entities in the pathway have been highlighted.

Let us denote with [SigH] and [SpaS] the concentration of the respective species, and with $\overline{[\text{SigH}]}$ and $\overline{[\text{SpaS}]}$ the corresponding averaged values throughout the whole environment. In the following, a modification of the model in [Hu *et al.*, 2004] is described.

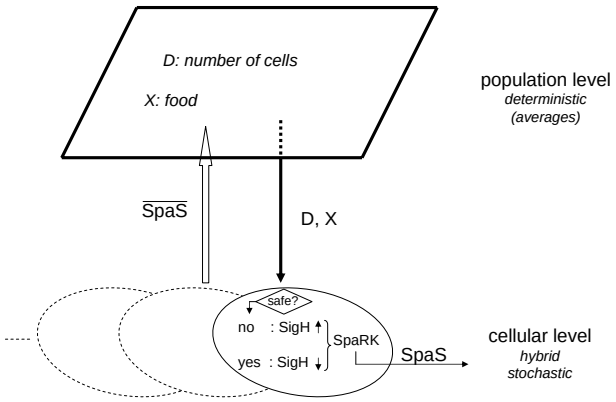

Figure 2.23: The decentralized structure of the model under study.

Let us start from the lower-level relations. In words, this part of the pathway depends on the dynamics of SigH, which is thought to be the actual initiator of the subtilin production mechanism. The dynamics of SigH depend on some external signals, the form of which we aim at unveiling. SigH influences the dynamics of SpaS[7] through a switching mechanism. This switching condition is encoded as a two-state Markov Chain: the higher the concentration level of SigH, the more likely it is that SpaS is produced, and hence that its level increases.

The level of the sigma factor SigH follows a controlled switching behavior:

$$\frac{d[SigH]}{dt} = \begin{cases} -\lambda_1[SigH] & \text{if production is OFF} \\ -\lambda_1[SigH] + k_3 & \text{if production is ON.} \end{cases}$$

In [Hu *et al.*, 2004] the conditions determining the ON/OFF status depended on a fixed, arbitrary threshold on the food level. This was suggested by the observed dependence of the production mechanism on some sort of *quorum sensing* [Tjalsma *et al.*, 2004]. In [Kouretas *et al.*, 2006b], instead, an identification scheme is employed to

[7]Notice again from Figure 2.22 that we skip the influence of the SpaRK cluster, because of the symmetry in the relation between SigH:SpaRK and that between SpaRK:SpaS.

find out parameters of the system, which can potentially represent the position of the switching thresholds for the system. The underlying hypothesis for this identification procedure to yield a meaningful outcome is that the thresholds have to be constant levels on any of the model coordinates. Notice further that, since all the models built so far are time-invariant, no threshold dependence on the growth phase of the species is taken into account [Stein, 2005]. The two mentioned approaches rule out the possibility of having a more complex dependence on the state space, which is indeed what this work aims at finding. Assuming fictitiously that the production can be "controlled," it is instructive to write out the above dynamical relations as follows:

$$\frac{d[SigH]}{dt} = -\lambda_1[SigH] + k_3 u. \tag{2.45}$$

The essential assumption is that the control has the form

$$u = f(D, X, [SigH], [SpaS]) : \mathcal{D} \to \{0, 1\}, \tag{2.46}$$

where $\mathcal{D} = [0, D_M] \times [0, X_M] \times \mathbb{R}^2_+ \subset \mathbb{R}^4$. This state feedback represents a general binary function of the variables of the model. Let us leave the actual shape of this control function to be determined—actually, we will argue that the synthesis of this control will be the outcome of the procedure this work is about to set up.
Next, the concentration of the protein SpaS depends on one of the two possible states of a switch S_1 as:

$$\frac{d[SpaS]}{dt} = \begin{cases} -\lambda_3[SpaS] & \text{if } S_1 \text{is OFF} \\ -\lambda_3[SpaS] + k_5 & \text{if } S_1 \text{is ON.} \end{cases} \tag{2.47}$$

$S_1 = \{OFF, ON\}$ is assumed to be a Markov Chain, whose transition probability matrix is:

$$P_1 = \begin{bmatrix} 1 - b_0 & b_0 \\ b_1 & 1 - b_1 \end{bmatrix}. \tag{2.48}$$

The coefficients b_0, b_1 depend directly on [SigH] according to [Hu *et al.*, 2004]

$$b_0([SigH]) = \frac{e^{-\Delta G_{rk}/RT}[SigH]}{1 + e^{-\Delta G_{rk}/RT}[SigH]}, \quad b_1([SigH]) = 1 - b_0([SigH]).$$

The quantity ΔG_{rk} represents the Gibbs free energy of the molecular configuration, R is a gas constant and T the environment temperature in Kelvin. This choice of transition probabilities makes S_1 reversible [Durrett, 2004], which helps in its analysis in that the steady state corresponds to the vector $[\pi_{OFF}, \pi_{ON}]^T = [b_1, b_0]^T$. Intuitively, as discussed above, SigH promotes the production of SpaS by increasing the likelihood of S_1 to be in the ON state.

At the higher level, the dynamics of the population and of the food level will be influenced by the amount of subtilin currently present in the environment. As the modeling choice was not to describe the dynamics of each single bacterial cell, we shall only account for the average production of subtilin, computed from the current antibiotic level for the single cell and depending from the environment in terms of the food and population levels.
The food level will directly depend on this average function, while the population level shall depend indirectly from it, through the value of the food.
The variation in the population level is modeled by a logistic equation as follows:

$$\frac{dD}{dt} = rD\left(1 - \frac{D}{D_\infty}\right). \tag{2.49}$$

It is a quadratic equation, with two equilibria. It is considered purposely for its clear asymptotics. The non-trivial equilibrium relation depends on the quantity D_∞, which is called the *carrying capacity*. Let us a priori define it to be equal to

$$D_\infty = \frac{X}{X_M} D_M,$$

where D_M and X_M represent the maxima for the population and the food levels in the environment. In other words, as intuitive, the steady state dynamically depends on the relative quantity of food in the environment. The food dynamics are taken to be:

$$\frac{dX}{dt} = -k_1 DX + k_2 \overline{[SpaS]}. \tag{2.50}$$

The above dynamical relation says that the food gets consumed at a rate proportional to its present level and the population density, while the average production of subtilin decreases this rate. This is due to the indirect negative influence of the presence of

the antibiotic on the population level. This average quantity is introduced to be the following:

$$\overline{[SpaS]} = \frac{D}{D_M}\left(1 - \frac{X}{X_M}\right)\frac{k_5}{\lambda_3}\bar{b}_0 h(X), \tag{2.51}$$

where $\bar{b}_0 = b_0(\overline{[\text{SigH}]})$ and $\overline{[\text{SigH}]} = \frac{k_3}{\lambda_1}$, and $h(X)$ is the Heaviside step function over the food variable, i.e. it is equal to 1 if $X > 0$, and to 0 if $X = 0$. The above relation stresses two separate influences. First, a dependence on the "competition" in the environment, given by first two terms. The average will be positively influenced by the relative population level (more bacteria induce a higher stress), whereas it will be negatively influenced by the amount of food (less food will increase the competition in the environment). Secondly, the average level of antibiotic depends on the steady-state dynamics for [SpaS] (fraction k_5/λ_3), that is on the production level for the single cell, which in turn depends on the steady state of SigH ($\pi_{ON} = b_0(k_3/\lambda_1)$, taken from the equations (2.45-2.47)).

From the dynamical relations in (2.49-2.50) for the higher-level variables and equation (2.51), the steady state of the population and food variables is going to be either $[D_{eq}, X_{eq}]^T = [0, \beta]^T, \beta \in [0, X_M]$, or $[D_{eq}, X_{eq}]^T = [\alpha D_M, \alpha X_M]^T$, where $\alpha \leq \frac{k_2 k_3 k_5}{k_2 k_3 k_5 + k_1 \lambda_1 \lambda_3 D_M X_M} \leq 1$. While the first equilibrium, as we discussed, is unstable, the second is stable for any combination of the model parameters.

The above set of dependencies and dynamical relations can be formally reframed as a SHS. In order to relate these dynamics to the framework in this chapter, from now on we shall work in discrete time, assuming that the above dynamical relations have been properly approximated, for instance with a simple first-order time discretization method. Notice that the probabilistic terms in the model are concentrated on the switching structure of S_1. The hybrid state space is made up of a discrete component θ (the state of S_1), and a continuous one (a vector in $\mathcal{D} \subset \mathbb{R}^4$). The hybrid state space will be denoted as $\mathcal{S} = \{OFF, ON\} \times \mathcal{D}$. The dynamics of SigH depend on a binary function that depends, according to equation (2.46), on a general feedback contribution f, function of the (continuous component of the) state space. Following [Hu *et al.*, 2004], it could for instance be assumed that f is piecewise-constant according to a proper spatial guard that depends on the food variable. Another discrete hybrid component would then be obtained.

However, in this work let us retain most generality: the objective will then be that of synthesizing this function according to certain optimality criteria that will be made explicit in the following. The shape of the outcome of this optimal control problem will then be fed back into the dynamics of SigH and will dictate whether it will present a threshold (as suggested in the literature) or, more generally, be non-linear function of the whole state space.

As in Definition 15, a control profile over a certain finite (discrete) time horizon $[0, N]$, i.e. a sequence of mappings $\mu = (\mu_0, \mu_1, \ldots, \mu_N)$ of the form in (2.46), shall be named a *strategy*, or a *policy*. The simple binary control space, denoted $\mathcal{U} = \{0, 1\}$, is in this case discrete and finite.

A solution of the above SHS model is a stochastic process with two components $\mathbf{s}(k) = (\theta(k), \mathbf{x}(k)), k \in [0, N]$ which, given an initial condition at time $k = 0$ (possibly sampled from an initial probability distribution), evolves in either of the two discrete modes until a jumping condition is verified (which reduces to sampling, along the evolution of the trajectory, from the inhomogeneous probability distribution of the MC S_1). Once a transition is triggered, the discrete state changes mode, while the continuous state remains unchanged (in the hybrid systems parlance, it is said the deterministic "reset" function is the identity), and the evolution continues from the unaltered conditions in the new mode. By construction, the solution of the model is a controlled stochastic process. Furthermore, given the structure of the policy and its sole dependence on the present state at each time step, once a strategy is selected the process is simply Markovian. Along with the general SHS interpretation, the model can be also thought as being a piecewise-deterministic Markov process, as in [Davis, 1993; Kouretas *et al.*, 2006b].

2.4.3 Survival Analysis as Probabilistic Safety Verification

The literature on antibiotic synthesis as a stress response for *B. subtilis* suggests that the production activation or de-activation follows some sort of "switching" profile [Stein *et al.*, 2002; Stein, 2005]. This observation has been thus far interpreted as the presence of specific spatial thresholds, that is precise values that are function either of the species concentration, or of the food or the population level, which are

characteristic to these switching behaviors. A research thrust has focused, assuming a special structure for these functions, on automatically identifying these thresholds from the data [Kouretas *et al.*, 2006b].

As motivated in the preceding section, in this work a rather different perspective is taken. The presence of the thresholds will not be a-priori postulated, but possibly obtained with respect to a certain survival property. That is, these switching conditions may be obtained as the outcome of a specific control synthesis problem. The use of the word "possibly" is stressed because while a binary function has been specifically chosen for the feedback control, an outcome with the shape of a threshold may not necessarily result from the synthesis problem. It is then key to make use of a proper survival function.

According to the claim discussed in the previous section, it is assumed that a species elicits or deactivates the "production pipeline" for the antibiotic with the main objective of maximizing its own survival likelihood. From an optimal control perspective, survival is encoded as an objective function for the single cell. Now, reinterpreting survival within the dynamic model introduced in section 2.4.2, it is possible to introduce certain safety regions within the state space that are associated with a specific survival status. In a deterministic setting, the survival status could be "alive" or "dead." Instead, in a less coarse stochastic modeling setup, we may better opt for defining a survival probability. Next, because a solution process of the SHS is dependent on a Markovian control, we aim at maximizing the chances of survival by synthesizing an optimal strategy. Given an educated choice for the survival function, let us claim that this strategy will turn out to be "of switched kind." In other words, let us argue that rather than associating the production activations/deactivations to spatial thresholds, as in the previous literature, these should be referred to certain *safety levels* for the species under study, once a proper survival criterion is selected.

Clearly, the above argument hinges on the relationship between survival and safety. The choice of selecting a proper survival function then translates into being able to formalize a proper definition of safety for the case under study. According to the interpretation for the present model, associating a safety region to a survival

condition can be schieved by stating that the single *B-subtilis* bacterium is safe if

$$[SpaS] > \overline{[SpaS]}. \tag{2.52}$$

Informally, if the subtilin production level of the single species under study is higher than the average subtilin present in the surrounding environment, then the species will be deemed to be safe. The opposite interpretation will hold for the complement of the claim. Intuitively, the relationship in (2.52) encodes a higher likelihood for the species to kill other bacteria, rather than being killed by their antibiotic. Equivalently, exploiting the expression in equation (2.51), let us define the *safe region* $\mathcal{A}$ to be the set of points

$$\mathcal{A} = \left\{ \boldsymbol{s} \in \mathcal{S} : [SpaS] > \frac{D}{D_M} \left(1 - \frac{X}{X_M} \right) \frac{k_5}{\lambda_3} \overline{b}_0 h(X) \right\}.$$

In general terms, in a stochastic setting a *safety analysis* problem consists in evaluating the probability that the state of the system remains outside a certain set deemed to be *unsafe* during a given time horizon, starting from some set of initial conditions. The objective of the study then becomes that of modeling antibiotic production via safety analysis in a SHS framework. Notice that the above survival objective, as well as the corresponding safety specification, do not include any energy cost as a possible penalty term. This is due to the model of choice, which does not include any food consumption term at the level of the single species. However notice that, as we are interested in focusing on the presence of thresholds in the production mechanism, rather than understanding its compete structure, this limitation plays a second role. In general, a performance criterion could be introduced in the survival analysis and retained within the safety interpretation by exploiting the ideas in section 2.2.7.

Having set up the above mathematical machinery, it is now possible to apply it to the problem under study. This can be done by leveraging the results in the first part of this chapter. The procedure outlined in Thm. 5 is practically implemented via a dynamic programming algorithm. This procedure, along with the optimal safety level $\mathrm{p}_s^{\mu^*}(\mathcal{A})$ associated with any point in the state space, yields also the optimal state-dependent, time-varying policies μ^*. This second output is really the focus of

our attention, as it shall represent the synthesized control functions for the relation (2.45), to be interpreted as the activation thresholds for the production of subtilin.

2.4.4 Numerical Results and Discussion

In this section let us report the outcomes of the simulations for the following experimental setup. The parameters have been chosen to be the following: $r = 0.8, k_1 = 2, k_2 = 4, k_3 = 2.5, k_5 = 0.8, \lambda_1 = 0.5, \lambda_3 = 0.2, \Delta G_{rk}/RT = 1.1$. The time horizon has been set to $N = 40$. In Fig. 2.24, the plots of the maximal probabilistic safe

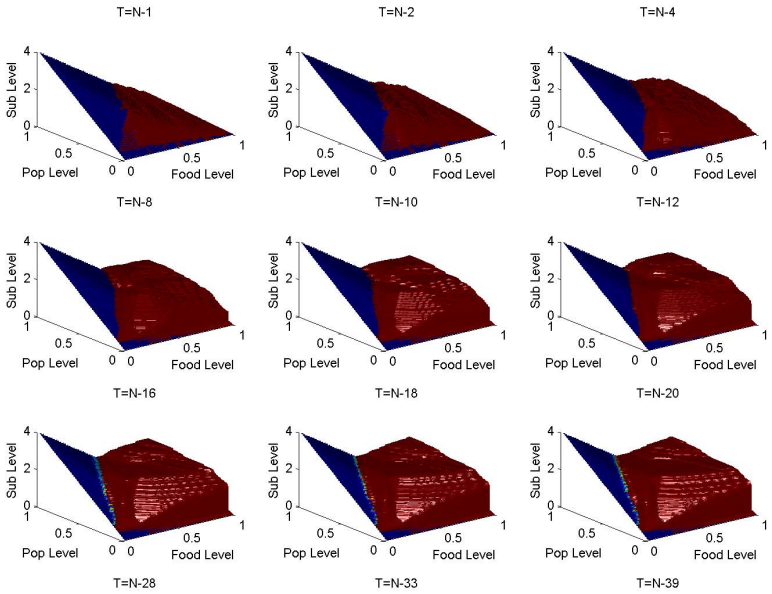

Figure 2.24: Maximal probabilistic safety level set corresponding to $\epsilon = 1$, backwards in time.

level sets, with safety level $\epsilon = 1$, are shown for different time samples. The colors have been added only to enhance the perspective and the height. Given the choice of the safety level ϵ, all the points in the plots above the surfaces are to be considered to be "almost surely" safe. Notice that, as expected, the safe set shrinks as we proceed backwards in time (this in fact translates to a longer safety requirement for the trajectories of the system). For the sake of visualization, we plotted the results corresponding to a fixed value of the sigma factor [SigH] = 1. All the plots refer to the discrete state being in the OFF mode. Fig. 2.25 represents, for a particular fixed

time sample, different maximal probabilistic safety level sets, each corresponding to a particular safety probability ϵ. For each level set, the corresponding safety region includes all the points above the set. The lower the probabilistic safety level, the larger the safety region. Finally, Fig. 2.26 represents pairs of plots referring to max-

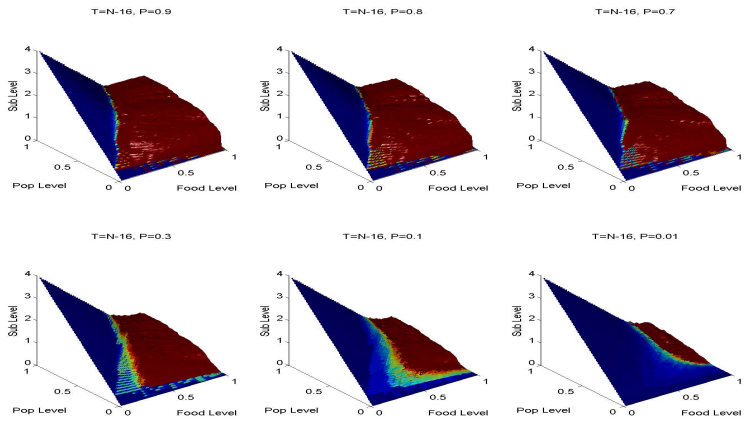

Figure 2.25: Maximal probabilistic safety level set corresponding to different safety levels (decreasing along the left-right and top-bottom direction).

imal probabilistic safe sets (for $\epsilon = 0.95$) and corresponding optimal actions. More precisely, the green plots (second and fourth row) represent the regions in the state space that are associated with a *switching action*, and are to be matched with the safety level sets plotted directly above them.

It is interesting to realize that the optimal control functions, associated with the activation thresholds, have a characteristic "onion layer" shape, varying along time. It can be thus argued that the optimal actions single out switching surfaces corresponding to certain safety levels. These surfaces have profiles that match the variation in safety probability (or, according to our interpretation, in survival likelihood, for the selected survival criterion) for the species, as appears by comparing the control plots with the safety level sets. In general, these surfaces are not hyper-rectangular, as the previous research efforts that sought to identify thresholds related to them assumed [Kouretas *et al.*, 2006a]. Instead, they are rather nonlinear functions of the state space, showing a manifest but non-explicit dependence with the change in safety level of the single species.

To make sure that the obtained outcome is indeed of value to the problem, similar

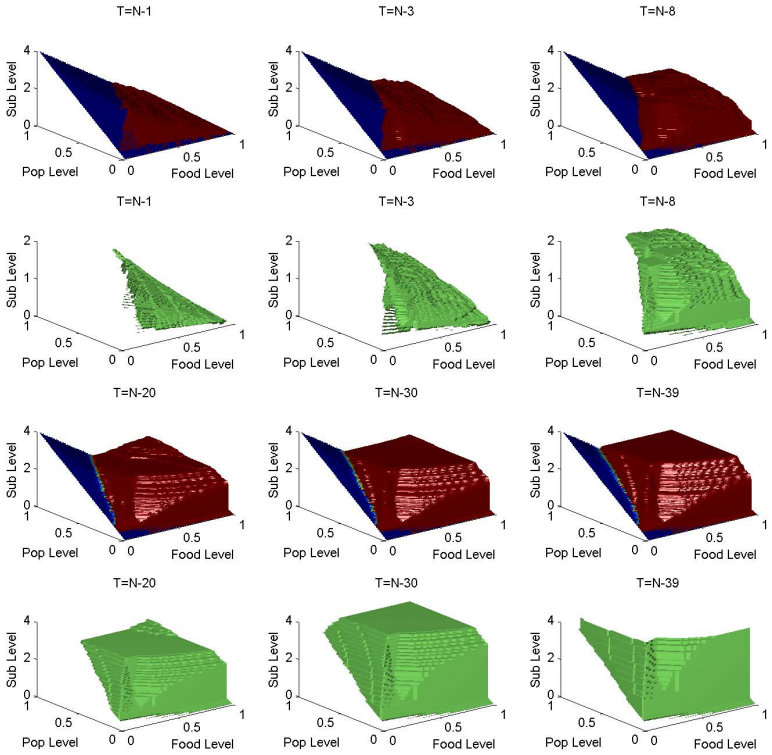

Figure 2.26: Maximal probabilistic safety level sets and optimal switching control, backwards in time.

simulations are implemented, which are instead based on a survival function depending on a simple feedback describing the competition coming from the environment. This is suggested by two ideas: first, the observation in the literature that the bacterium adjusts its behavior according to some sort of *quorum sensing* [Tjalsma *et al.*, 2004]; second, the tentative of the former modeling efforts to define thresholds which depend only on food level and population density [Hu *et al.*, 2004]. More precisely, the set of safe points can be defined to be:

$$\mathcal{A} = \left\{ \boldsymbol{s} \in \mathcal{S} : \frac{D}{D_M} \left(1 - \frac{X}{X_M} \right) > \texttt{thresh} \right\}.$$

Notice that the subtilin level influences the food level, and indirectly also the population level. In the above, we have assumed a value $\texttt{thresh} = 0.3$. The outcomes of the simulations (to be interpreted as the above ones in figure 2.26), shown in figure 2.27, do not appear to yield any threshold behavior for the production mechanism. Similarly, some experiments performed by assuming a safe set which is defined on the food coordinate (as in [Hu *et al.*, 2004]) were run. However, the outcomes (see figure 2.27) did not show any particular "threshold," thus suggesting that such boundary may not be the actual discriminant for the subtilin production mechanism.

Let us remark that the argument developed in this section leads to conclude that the observed production mechanism shows a switched feature, which is associated with a condition that, rather than being referred to some value for the external coordinates of the system, has instead to be interpreted in terms of a particular level of the survival probability for the species.

Clearly, this approach subsumes the ability to define a specific survival function, or a corresponding safety set, for the system under study. As already discussed above, in general it may be argued that some energy-related terms ought to be included in a cost with the shape of (2.52). A wrong choice for this function may lead to misleading conclusions, as described above.

From a biological perspective, it is necessary to understand what are the exogenous "signals" that the species is able to sense, and which can presumably build up the survival function. Our choice of an average feedback may not be valid in general. Furthermore, the present study assumes that each cell optimizes its own survival, thus

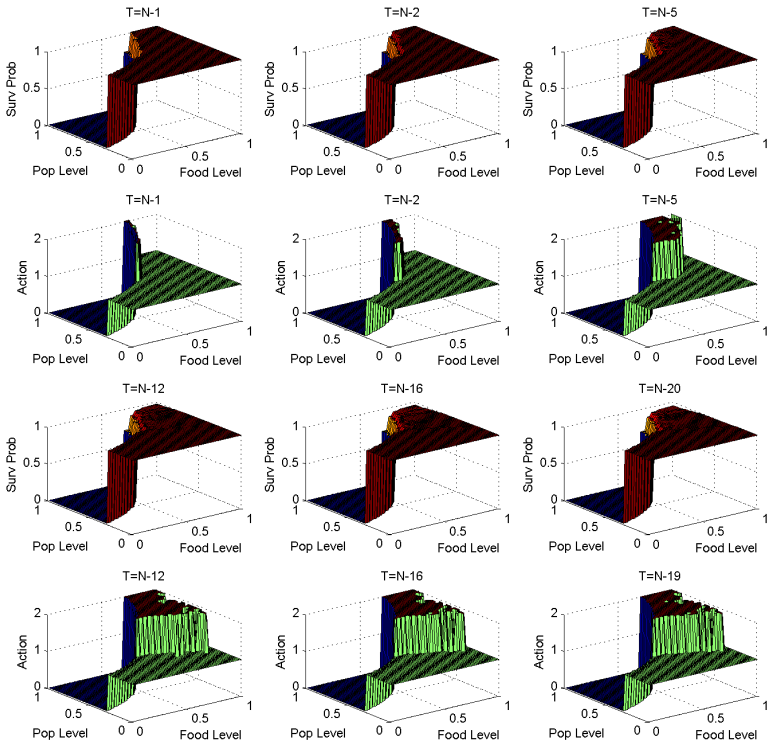

Figure 2.27: Maximal probabilistic safety level sets and optimal control, backwards in time, based on a fitness function encoding competition in the environment.

ruling out any possible cooperative, or emerging global behavior. As we discussed in the introduction to this section, this idea may be contended: while this arguably makes sense for the current study, the reader is urged to ponder over the general meaning of the "unit of selection."

As an extension to the present work, it would be instructive to understand what is the *critical* safety level that corresponds to the activation of the production of subtilin. Our computations hint at showing that, for the particular case under study, the level is a safety probability value close to one. More accurate testing needs to be performed on this aspect.

2.5 Further Theoretical Investigations and Computational Studies

There are a number of problems that directly extend those already tackled in this section, as well as other issues that stem from the work presented in this dissertation. Let us leave the following points as possible avenues for future research.

Sufficiency of Markov Policies. Much like other stochastic optimal control problems appeared in the literature, in this theoretical setup it looks likely that Markov policies are optimal within the set of general feedback policies. This statement is fairly straightforward to prove for additive cost functions, under weak assumptions on the entities into play [Kumar and Varaiya, 1986; Bertsekas and Shreve, 1996]. However, at this time the formal proof of this statement for either the max or the multiplicative cost functions used in this work appears to be elusive.

Game-theoretical framework. Modeling adversarial uncertainty and synthesizing control strategies for this problem setup has an intrinsic interest, as well as a clear and direct applicability to a number of important control problems. Results originating from this line of research would also connect with related efforts in the SHS community [Mitchell and Templeton, 2005]. Please refer to section 2.2.6 for some preliminary details.

Randomized Strategies. It is certainly of interest to extend the developed framework to controls that, while being dependent on the state space, are not necessarily deterministic.
From a theoretical perspective, the introduction of randomized policies requires special care. For instance, the non-anticipativity of such signals may be assumed, that is independence between the filtration they induce in the past, and the filtration resulting from the other probabilistic terms in the future dynamics. The work in [Ghosh *et al.*, 1992; Ghosh *et al.*, 1997] contains a number of details on how to properly handle such instance, and solve related optimal control problems.
The benefits of this approach are its generality and its possible connection with in-

teresting engineering applications. We have mentioned in section 2.2.6, for instance, that this would model a key feature of *fault-tolerant* control. The tradeoff on this point is that the solution of a control synthesis problem for randomized problems may be computationally heavier. This suggests that the computational aspect may be a limiting factor for the proposed approach.

Computational Improvements. Along the above lines, which emphasize the dimensionality limitations of the proposed techniques, the study of more efficient approaches to solve the DP algorithm on a hybrid state space is of primary focus. The literature suggests some methods to attack this problem. As introduced in section 2.3.4, one technique exploits some decentralization in the structure of the dynamics in order to distribute the computations: HS models naturally yield themselves to this distributed approach according to the topology of the underlying graph, consisting of the modes and the edges of the HS. However, notice that this does not generally reduce the "continuous" part of the dimension. In [Bertsekas, 1982], an approach to asynchronously perform in parallel the computations – with proven convergence – is suggested.

A second, more recent approach suggests to solve large-scale Markov Decision Processes (MDPs) by approximating the optimal value function by a linear combination of basis functions and finding the associated optimal weights by linear programming [Kveton *et al.*, 2006].

Another idea is that of exploiting an automated control synthesis toolbox to ease the solution of the DP scheme. The work in [Baotic *et al.*, 2003; Christophersen *et al.*, 2005] bridges the solution of a general DP program with that of a model predictive control (MPC) one. The latter optimization setup can be solved by the multi-parametric toolbox (MPT), for a number of model instances which depend on the dynamical structure (for instance, in the piece-wise affine case) and the cost functions (for example, in the linear case). The author is currently investigating and experimenting with these methods to achieve computationally attractive performances for the proposed schemes, to be further tested on the benchmark in [Fehnker and Ivančić, 2004] and compared to the other techniques in the literature.

Embedding Performance. As it has been argued in section 2.2.7, the author is investigating ways to "hide" some performance specifications in the reachability/safety framework under study. From a heuristic point of view, and embracing the state-augmentation technique proposed in section 2.2.7, a number of approaches can be implemented. According to these approaches, certain forms of discounting can be used, which allow to express a safety specification problem as a norm minimization one. While practically functional, these approaches often are not prone to be formally expressed and investigated. Furthermore, they often tarnish the underlying probabilistic interpretation of the output of the optimization problem (which may be still recovered a-posteriori).

More generally, embedding further performance specifications in the control design procedure, while guaranteeing safety, is a well studied problem. An interesting approach which could also be inspiring for this research is that proposed in [Lygeros *et al.*, 1999], addressing a multi-objective control problem with requirements ranking for deterministic hybrid systems. This can have important applicative outcomes, and yield interesting new interpretations of the problem under study.

Connections with Continuous Time. An important issue to be addressed is how to extend the proposed methodology for probabilistic reachability analysis and design to a continuous time stochastic hybrid setting. In section 2.2.9 we have connected and compared the presented work with other approaches. We have further strengthened this connection by discussing the limiting equivalence of approximated models in the SHS literature (see chapter 1). Upon extending such approximation procedure to the controlled case, we would like to run fair comparisons of the two methods under the common benchmark in [Fehnker and Ivančić, 2004].

Discretization of the Control Space. As mentioned in section 2.3.3, the extension of the discretization approach to the control space has to be formally pinned down. However, it appears to be straightforward from that obtained for the state space, which has been presented in section 2.2.8.1.

The next two subjects refer to the theory developed in chapter 1.

Stochastic Hybrid Systems with no spatial guards. The author is interested in further exploiting the generator approach in section 1.3.3 and the formal and computational simplicity of SHS models with no spatial conditions. This may be of interest in investigating other problems that are not strictly related to the probabilistic reachability and safety ones.

A Time Discretization Scheme with Stronger Convergence. In connection with the theory developed in section 1.4, the author is working on an extension of the results towards proving some stronger notion of convergence. The main reference work for these results is [Krystul and Bagchi, 2004], which heavily hinges on results formerly shown in [Gobet, 2000]. More restricted results are proven with similar techniques in [Yuan and Mao, 2004], where the framework of switching diffusions is taken in consideration. Furthermore, higher-order discretization schemes can be obtained as extensions of the presented one [Kloeden and Platen, 1992].

✠

Symbols, Notations, and Acronyms

Let us recapitulate here the most common notations, symbols and conventions from the literature, that are used throughout the text. We exclude from this list recent or novel concepts formally introduced in the thesis. These quantities are also formally introduced or defined when first used.

- $\mathbb{N}$ denotes the natural numbers, $\mathbb{Z}$ is the set of integers, $\mathbb{R}$ are the real numbers, $\mathbb{R}^+$ the positive real numbers, $\mathbb{R}_0^+$ the nonnegative reals.
- A trajectory, an execution, or a stochastic process is denoted with bold fonts, whereas values that this function may assume are denoted with regular fonts (even if they may be vectors).
- When taking the expectation of (a function of) a Markov process, let us use a subscript to denote the initial distribution of the process, or the initial value of the process: $E_{x_0}\left[\mathbf{x}(t)\right] = E\left[\mathbf{x}(t)|\mathbf{x}(0) = x_0\right]$.
- The probability that a certain event is true for a stochastic process, conditioned on its initial condition, will be denoted with a subscript, with the initial condition: $\mathcal{P}_{x_0}\left(\mathbf{x}(t) \in A\right) = \mathcal{P}\left(\mathbf{x}(t) \in A|\mathbf{x}(0) = x_0\right)$.
- Ω is the underlying event space over which let us define the probability space assigned to a stochastic model, see page 21.
- $\mathscr{F}$ is the natural filtration.
- $\mathcal{P}$ is the probability assigned to events on a sigma-field, or to trajectories of a system.

- $\mathcal{N}(m, \sigma)$ denotes a normal probability distribution with mean m and variance σ^2. If the two parameters are an n-dimensional vector and a $x \times n$ symmetric, positive semidefinite matrix, then $\mathcal{N}(m, \sigma)$ is a multi-dimensional distribution.

- $\mathscr{U}([a, b])$ denotes a uniform probability distribution over the interval $[a, b], a < b \in \mathbb{R}^n, 1 \leq n < \infty$.

- Given a E-valued random variable X, the function $\texttt{dist}_X : E \to [0, 1]$ denotes its probability distribution. Further, let us say that $\texttt{dist}_X$ belongs to $\mathscr{P}(E)$, the space of probability distributions over E.

- Unless otherwise stated, $\| \cdot \|$ is the supremum norm.

- Given a point z and a set A in a normed space (like the Euclidean space endowed with its natural norm), let $d(z, A) = \inf_{y \in A} ||z - y||$ denote the distance between z and A. Similarly if A is made up by a single point.

- Given two numbers $a, b, a \vee b = \max\{a, b\}$.

- Given a sequence $(a_n)_{n \geq 1}$ and an a, $\lim_n^\star a_n = a \Leftrightarrow \lim_{n \to \infty} a_n = a$ and $(\vee_n \|a_n\|) \vee \|a\| < \infty$.

- The function `card`, when applied to a general discrete set, yields the cardinality of that set, that is the number its elements.

- `Leb` denotes the Lebesgue measure over the specified space.

- $\delta(n), n \in \mathbb{Z}$ denotes the Kronecker delta extended to the integers.

- The empty set is denoted as $\varnothing$.

- Given a set A, its complement is denoted by A^c, or $\bar{A}$.

- Given two homogeneous sets A, B, their difference is defined as $A \backslash B = \{x : x \in A \wedge x \notin B\}$.

- Given a pair (usually referred to an edge in a HS model) $e = (e_1, e_2)$, we have introduced $\mathfrak{s}(e) = e_1$ and $\mathfrak{t}(e) = e_2$.

- Given a vector x of dimension n, for instance $x \in \mathbb{R}^n$, $\texttt{diag}(x)$ denotes the $n \times n$ square diagonal matrix, with elements taken orderly from vector x. The n-dimensional identity matrix, made up of unity terms along the diagonal and null terms elsewhere, is denoted by I_n.

- The identity function maps points in a space into themselves, for instance, $\forall x \in \mathbb{R}^n, \texttt{id}(x) = x$.

- Consider a function $f : A \to A'$. If $f \in B(A)$, then f is Borel-measurable in its domain of definition; if $f \in C(A) = C^0(A)$, then f is continuous in its domain of definition A; if $f \in C_b(A)$, then f is continuous and bounded in its domain of definition; if $f \in C_b^n(A), n \geq 1$, then f is n-times continuously differentiable and bounded; if $f \in \bar{C}_b^n(A), n > 1$, then f is n-times continuously differentiable and uniformly bounded.

- A function $f : A \subseteq \mathbb{R}^+ \to \mathbb{R}$ is said to be *càdlàg* if, for any $t \in A, \lim_{s \downarrow t} f(s) = f(t)$, and if $\lim_{s \uparrow t} f(s)$ exists.

- *a.s.* = *almost surely*. When a property holds a.s. on a probability space $(\Omega, \mathscr{F}, P)$, it does so for every events in Ω, except possibly a set A such that $P(A) = 0$.

- OC = Optimal Control

- HS = Hybrid System

- SHS = Stochastic Hybrid System

- GSHS = General Stochastic Hybrid System. Throughout the thesis, the term GSHS is often associated with a SHS model with a non-trivial guard set, while SHS to a model with no spatial conditions.

- CTSHS = Continuous-Time Stochastic Hybrid System

- DTSHS = Discrete-Time Stochastic Hybrid System

- DP = Dynamic Programming
- CCC = Compact Containment Condition
- ODE = Ordinary Differential/Difference Equation
- SDE = Stochastic Differential Equation
- PDE = Partial Differential Equation
- PDMP = Piecewise-Deterministic Markov Process

✠

Bibliography

[Abate and Tiwari, 2006] A. Abate and A. Tiwari. Box invariance of hybrid and switched systems. In *Proceedings of the 2nd IFAC Conference on Analysis and Design of Hybrid Systems*, Alghero, IT, June 2006.

[Abate *et al.*, 2004] A. Abate, L. Shi, S. Simic, and S. Sastry. A stability criterion for stochastic hybrid systems. In *Proceedings of the 16th International Symposium on Mathematical Theory of Networks and Systems*, Leuwen, BG, July 2004.

[Abate *et al.*, 2005] A. Abate, A. Ames, and S. Sastry. A stochastic approximation for hybrid systems. In *Proceedings of the 24th American Control Conference*, Portland, OR, June 2005.

[Abate *et al.*, 2006a] A. Abate, A. Ames, and S. Sastry. A-priori detection of zeno behavior in communication networks modeled as hybrid systems. In *Proceedings of the 25th American Control Conference*, pages 1564–1569, Minneapolis, MN, June 2006.

[Abate *et al.*, 2006b] Alessandro Abate, Alessandro D'Innocenzo, Giordano Pola, Maria Domenica Di Benedetto, and Shankar Sastry. The concept of deadlock and livelock in hybrid control systems. Technical Report UCB/EECS-2006-181, EECS Department, University of California, Berkeley, December 17 2006.

[Alur and Dill, 1994] R. Alur and D. Dill. A theory of timed automata. *Theoretical Computer Science*, 126:183–235, 1994.

[Alur *et al.*, 1993] R. Alur, C. Courcoubetis, T.A. Henzinger, and P. Ho. Hybrid automaton: an algorithmic approach to the specification and verification of hybrid systems. In R.L. Grossman, A. Nerode, A.P. Ravn, and H. Rishel, editors, *Hybrid Systems*, LNCS 736, pages 209–229. Springer Verlag, 1993.

[Alur *et al.*, 2000] R. Alur, T. Henzinger, G. Lafferriere, and G.J. Pappas. Discrete abstractions of hybrid systems. *Proceedings of the IEEE*, 88(2):971–984, 2000.

[Alur *et al.*, 2001] R. Alur, C. Belta, F. Ivančić, V. Kumar, M. Mintz, G.J. Pappas, H. Rubin, and J. Schug. Hybrid modeling and simulation of biomolecular networks. In M.D. Di Benedetto and A. Sangiovanni-Vincentelli, editors, *Hybrid Systems: Computation and Control*, LNCIS 2034, pages 19–32. Springer Verlag, 2001.

[Ames and Sastry, 2005a] A. Ames and S. Sastry. A homology theory for hybrid systems: Hybrid homology. In M. Morari and L. Thiele, editors, *Hybrid Systems: Computation and Control*, LNCIS 3414, pages 86–102. Springer Verlag, 2005.

[Ames and Sastry, 2005b] A. D. Ames and S. Sastry. Characterization of Zeno behavior in hybrid systems using homological methods. In *24th American Control Conference*, Portland, OR, 2005.

[Ames *et al.*, 2005] A. D. Ames, A. Abate, and S. Sastry. Sufficient conditions for the existence of Zeno behavior. In *44th IEEE Conference on Decision and Control and European Control Conference*, Seville, SP, 2005.

[Ames *et al.*, 2007] A. D. Ames, A. Abate, and S. Sastry. Sufficient conditions for the existence of Zeno behavior in nonlinear hybrid systems via constant approximations. In *46th IEEE Conference on Decision and Control and European Control Conference*, New Orleans, LA, 2007.

[Anai and V.Weispfenning, 2001] H. Anai and V.Weispfenning. Reach set computations using real quantifier elimination. In M. D. Di Benedetto and A.L. Sangiovanni-Vincentelli, editors, *Hybrid Systems: Computation and Control*, volume 2034 of *Lecture Notes in Computer Science*, pages 103–117. Springer Verlag, 2001.

[Arens, 1949] Richard Arens. Approximations in, and representation of, certain banach algebras. *American Journal of Mathematics*, 71:763–790, 1949.

[Arnold, 1992] Ludwig Arnold. *Stochastic Differential Equations: Theory and Applications*. Krieger, 1992.

[Asarin *et al.*, 2000] E. Asarin, O. Bournez, T. Dang, and O. Maler. Approximate reachability analysis of piecewise linear dynamical systems. In N. Lynch and B. Krogh, editors, *Hybrid Systems: Computation and Control*, Lecture Notes in Computer Science 1790, pages 21–31. Springer Verlag, 2000.

[Asarin *et al.*, 2003] E. Asarin, T. Dang, and A. Girard. Reachability analysis of nonlinear systems using conservative approximation. In O. Maler and A. Pnueli, editors, *Hybrid Systems: Computation and Control*, Lecture Notes in Computer Science 2623, pages 20–35. Springer Verlag, 2003.

[Balluchi *et al.*, 2000] A. Balluchi, L. Benvenuti, M. D. Di Benedetto, G. M. Miconi, U. Pozzi, T. Villa, H. Wong-Toi, and A. L. Sangiovanni-Vincentelli. Maximal safe set computation for idle speed control of an automotive engine. In N. Lynch and B. Krogh, editors, *Hybrid Systems: Computation and Control*, Lecture Notes in Computer Science 1790, pages 32–44. Springer Verlag, 2000.

[Banerjee and Hansen, 1988] S. Banerjee and J. Hansen. Structure and expression of a gene encoding the precursor of subtilin, a small proteic antibiotic. *Journal of Biological Chemistry*, 263-19:9508–9514, 1988.

[Baotic *et al.*, 2003] M. Baotic, F.J. Christophersen, and M. Morari. A new algorithm for constrained finite time optimal control of hybrid systems with a linear performance index. In *In the Proceedings of the European Control Conference*, Cambridge, UK, 2003.

[Batina *et al.*, 2002] I. Batina, A.A. Stoorvogel, and S. Weiland. Optimal control of linear, stochastic systems with state and input constraints. In *Proceedings of the 41st Conference of Decision and Control*, pages 1564–1569, Las Vegas, NV, 2002.

[Bayen *et al.*, 2003] A. Bayen, S. Santhanam, I. Mitchell, and C. Tomlin. A differential game formulation of alert levels in ETMS data for high altitude traffic. In *Proceedings of the AIAA Congerence on Guidance, Navigation, and Control*, Austin, TX, 2003.

[Belta *et al.*, 2004] C. Belta, P. Finin, L. Habets, A. Halász, M. Imielinski, V. Kumar, and H. Rubin. Understanding the bacterial stringent response using reachability analysis of hybrid systems. In R. Alur and G. J. Pappas, editors, *Hybrid Systems: Computation and Control*, Lecture Notes in Computer Science 2993, pages 111–126. Springer Verlag, 2004.

[Bertsekas and Shreve, 1996] D. P. Bertsekas and S. E. Shreve. *Stochastic optimal control: the discrete-time case*. Athena Scientific, 1996.

[Bertsekas and Tsitsiklis, 1996] D. P. Bertsekas and J. N. Tsitsiklis. *Neuro-Dynamic Programming*. Athena Scientific, 1996.

[Bertsekas, 1971] D.P. Bertsekas. *Control of Uncertain Systems with a Set-Membership Description of the Uncertainty*. PhD thesis, Department of Electrical Engineering, MIT, 1971.

[Bertsekas, 1972] D. P. Bertsekas. Infinite-time reachability of state-space regions using feedback control. *IEEE Transactions on Automatic Control*, AC-17(5):604–613, 1972.

[Bertsekas, 1975] D. P. Bertsekas. Convergence of discretization procedures in dynamic programming. *IEEE Transactions on Automatic Control*, AC-20(3):415–419, 1975.

[Bertsekas, 1982] D. P. Bertsekas. Distributed dynamic programming. *IEEE Transactions on Automatic Control*, AC-27(3):610–616, 1982.

[Bhatt and Karandikar, 1993] A.G. Bhatt and R.L. Karandikar. Weak convergence to a markov process: the martingale approach. *Probability Theory and Related Fields*, 96:335–351, 1993.

[Billingsley, 1995] Patrick Billingsley. *Probability and Measure - Third Edition*. Wiley, 1995.

[Blake *et al.*, 2003] W.J. Blake, M. Kaem, C.R. Cantor, and J.J. Collins. Noise in eukaryotic gene expression. *Nature*, 422:633–637, 2003.

[Blom and Bloem, 2004] H.A.P. Blom and E.A. Bloem. Particle filtering for stochastic hybrid systems. In *Proceedings of the 43rd IEEE Conference on Decision and Control*, Paradise Island, Bahamas, Dec 2004.

[Blom and Lygeros, 2006] H.A.P. Blom and J. Lygeros. *Stochastic Hybrid Systems: Theory and Safety Critical Applications*. LNCIS, vol. 337, Springer Verlag, 2006.

[Blom, 2003] H. Blom. Stochastic hybrid processes with hybrid jumps. In *International Federation of Automatic Control Conference on Analysis and Design of Hybrid Systems*, Saint-Malo, France, 2003.

[Bolzern *et al.*, 2006a] P. Bolzern, P. Colaneri, and G. De Nicolao. Almost sure stability of stochastic linear systems with ergodic parameters: an average contractivity criterion. In *Proceedings of the 45th IEEE Conference on Decision and Control*, San Diego, CA, Dec 2006.

[Bolzern *et al.*, 2006b] P. Bolzern, P. Colaneri, and G. De Nicolao. On almost sure stability of continuous-time markov jump linear systems. *Automatica*, 42(6):983–988, 2006.

[Borkar *et al.*, 1999] V. S. Borkar, M. K. Ghosh, and P. Sahay. Optimal control of a stochastic hybrid system with discounted cost. *Journal of Optimization Theory and Applications*, 101(3):557–580, June 1999.

[Borkar, 1989] V. S. Borkar. *Optimal Control of Diffusion Processes*. Pittman Research Notes in Mathematical Series. Longman Scientific & Technical, 1989.

[Borkar, 1995] Vivek Borkar. *Probability Theory - An Advanced Course*. Springer Verlag, 1995.

[Branicky, 1994] M.S. Branicky. Stability of switched and hybrid systems. In *Proceedings of the 33rd IEEE Conference on Decision and Control*, Lake Buena Vista, FL, Dec 1994.

[Branicky, 1995] M.S. Branicky. *Studies in Hybrid Systems: Modeling, Analysis, and Control*. PhD thesis, Electrical Engineering and Computer Science, MIT, June 1995.

[Brockett, 1993] R. Brockett. Hybrid models for motion control systems. In H. Trentelman and J. Willems, editors, *Perspectives in Control*. Birkhauser, 1993.

[Bujorianu and Lygeros, 2003] M. L. Bujorianu and J. Lygeros. Reachability questions in piecewise deterministic makov processes. In O. Maler and A. Pnueli, editors, *Hybrid Systems: Computation and Control*, Lecture Notes in Computer Science 2623, pages 126–140. Springer Verlag, 2003.

[Bujorianu and Lygeros, 2004a] M.L. Bujorianu and J. Lygeros. General stochastic hybrid systems. In *Proceeding of 12th Mediterranean Conference on Control and Automation*, Aydin, Turkey, 2004.

[Bujorianu and Lygeros, 2004b] M.L. Bujorianu and J. Lygeros. General stochastic hybrid systems: Modelling and optimal control. In *Proceedings of the 43th IEEE Conference on Decision and Control*, Atlantis, Paradise Island, Bahamas, 2004.

[Bujorianu and Lygeros, 2004c] M.L. Bujorianu and J. Lygeros. Theoretical foundations of general stochastic hybrid processes. In *16th International Symposium on Mathematical Theory of Networks and Systems*, Leuven, Belgium, 2004.

[Bujorianu and Lygeros, 2006] M. L. Bujorianu and J. Lygeros. Toward a general theory of stochastic hybrid systems. In H.A.P. Blom and J. Lygeros, editors, *Stochastic Hybrid Systems*, LNCIS 337, pages 3–30. Springer Verlag, 2006.

[Bujorianu, 2004] M. L. Bujorianu. Extended stochastic hybrid systems and their reachability problem. In R. Alur and G. Pappas, editors, *Hybrid Systems: Computation and Control*, Lecture Notes in Computer Science 2993, pages 234–249. Springer Verlag, 2004.

[Chatterjee and Liberzon, 2004] D. Chatterjee and D. Liberzon. On stability of stochastic switched systems. In *Proceedings of the 43rd IEEE Conference on Decision and Control*, Paradise Island, Bahamas, Dec 2004.

[Chatterjee and Liberzon, 2006] D. Chatterjee and D. Liberzon. Stability analysis of deterministic and stochastic switched systems via a comparison principle and multiple lyapunov functions. *SIAM Journal on Control and Optimization*, 45(1):174–206, 2006.

[Cho *et al.*, 2005] K.H. Cho, K.H. Johansson, and O. Wolkenhauer. A hybrid systems framework for cellular processes. *BioSystems*, 80:273–282, 2005.

[Christophersen *et al.*, 2005] F.J. Christophersen, M. Baotic, and M. Morari. Optimal control of piecewise affine systems: A dynamic programming approach. In T. Meurer et al., editor, *Control and Observer Design*, LNCIS 322, pages 183–198. Springer Verlag, 2005.

[Chutinan and Krogh, 1998] A. Chutinan and B. Krogh. Computing polyhedral approximations to flow pipes for dynamic systems. In *Proceedings of the* 37^{th} *IEEE Conference on Decision and Control, Tampa, FL*, pages 2089–2094, December 1998.

[Cont and Tankov, 2004] Rama Cont and Peter Tankov. *Financial Modelling with Jump Proceses*. Chapman and Hall/CRC Press, 2004.

[Davis, 1984] M. H. A. Davis. Piecewise deterministic markov processes: a general class of nondiffusion stochastic models. *Journal of Royal Statistical Society series B*, 46:353–388, 1984.

[Davis, 1993] M. H. A. Davis. *Markov Models and Optimization*. Chapman & Hall/CRC Press, London, 1993.

[De Jong, 2002] H. De Jong. Modeling and simulation of genetic regulatory systems: A literature review. *Journal of Computational Biology*, 9(1):67–103, 2002.

[Digailova and Kurzhanski, 2004] I.A. Digailova and A.B. Kurzhanski. Attainability problems under stochastic perturbations. *Differential Equations*, 40(11):1573–1578, 2004.

[Digailova and Kurzhanski, 2005] I.A. Digailova and A.B. Kurzhanski. Reachability analysis under control-dependent stochastic noise. In *Proceedings of the 16th IFAC World Congress*, Prague, CZ, 2005.

[Duprè, 1987] J. Duprè. *The Latest on the Best: Essays on Evolution and Optimality*. MIT Press, Cambridge, MA, 1987.

[Durrett, 2004] Richard Durrett. *Probability: Theory and Examples - Third Edition*. Duxbury, 2004.

[Entian and de Vos, 1996] K. Entian and W. de Vos. Genetics of subtilin and nisin biosyntheses – biosynthesis of lantibiotics. *Antonie van Leeuwenhoek*, 69-2:109–117, 1996.

[Esposito *et al.*, 2001] J.M. Esposito, V. Kumar, and G.J. Pappas. Accurate event detection for simulating hybrid systems. In M.D. Di Benedetto and A. Sangiovanni-Vincentelli, editors, *Hybrid Systems: Computation and Control*, LNCIS 2034, pages 204–217. Springer Verlag, 2001.

[Ethier and Kurtz, 1986] Stewart N. Ethier and Thomas G. Kurtz. *Markov Processes, Characterization and Convergence*. Wiley, 1986.

[Evans and Souganidis, 1984] L. C. Evans and P. E. Souganidis. Differential games and representation formilas for solutions of Hamilton-Jacobi-Issacs equations. *Indiana University Mathematics Journal*, 33:773–797, 1984.

[Fehnker and Ivančić, 2004] A. Fehnker and F. Ivančić. Benchmarks for hybrid systems verifications. In R. Alur and G.J. Pappas, editors, *Hybrid Systems: Computation and Control*, Lecture Notes in Computer Science 2993, pages 326–341. Springer Verlag, 2004.

[Fleming and Souganidis, 1989] W. H. Fleming and P. E. Souganidis. On the existence of value functions of two-player, zero-sum stochastic differential games. *Indiana University Mathematics Journal*, 38:293–313, 1989.

[Geromel and Colaneri, 2006] J.C. Geromel and P. Colaneri. Stability and stabilization of discrete-time switched systems. *International Journal of Control*, 79(7):719–728, 2006.

[Ghosh and Bagchi, 2004] Mrinal Ghosh and Arunabha Bagchi. Modeling stochastic hybrid systems. In *System Modeling and Optimization. Proceedings of the 21st IFIP TC7 Conference.*, pages 269–279, Sophia Antipolis, FR, 2004.

[Ghosh and Marcus, 1995] M. K. Ghosh and S. I. Marcus. Stochastic differential games with multiple modes. Technical Report TR 1995-68, The Institute for Systems Research, University of Maryland, 1995.

[Ghosh *et al.*, 1992] M. K. Ghosh, A. Araposthasis, and S. I. Marcus. Optimal control of switching diffusions with application to flexible manufacturing systems. *SIAM Journal of Control and Optimization*, 30(6):1–23, November 1992.

[Ghosh *et al.*, 1997] M. K. Ghosh, A. Araposthasis, and S. I. Marcus. Ergodic control of switching diffusions. *SIAM Journal of Control and Optimization*, 35(6):1952–1988, November 1997.

[Gillespie, 1976] D.T. Gillespie. A general method for numerically simulating the stochastic time evolution of coupled chemical reactions. *Computational Physics*, 22:403–434, 1976.

[Gillespie, 1977] D.T. Gillespie. Exact stochastic simulation of coupled chemical reactions. *Physical Chemistry*, 81(25):2340–2361, 1977.

[Girard *et al.*, 2006] A. Girard, A. Julius, , and G. J. Pappas. Approximate simulation relations for hybrid systems. In *International Federation of Automatic Control Conference on Analysis and Design of Hybrid Systems*, Alghero, IT, 2006.

[Girard, 2005] A. Girard. Reachability of uncertain linear systems using zonotopes. In M. Morari and L. Thiele, editors, *Hybrid Systems: Computation and Control*, Lecture Notes in Computer Science 3414, pages 291–305. Springer Verlag, 2005.

[Glasserman and Merener, 2003] Paul Glasserman and Nicolas Merener. Numerical solution of jump-diffusion LIBOR market models. *Finance and Stochastics*, 7:1–27, 2003.

[Glasserman and Merener, 2004] Paul Glasserman and Nicolas Merener. Convergence of a discretization scheme for jump-diffusion processes with state-dependent intensities. *Proc. Royal Society*, 460:111–127, 2004.

[Glover and Lygeros, 2004] W. Glover and J. Lygeros. A stochastic hybrid model for air traffic control simulation. In R. Alur and G. J. Pappas, editors, *Hybrid Systems: Computation and Control*, Lecture Notes in Computer Science 2993, pages 372–386. Springer Verlag, 2004.

[Gobet, 2000] Emmanuel Gobet. Weak approximation of killed diffusion using Euler scheme. *Stochastic Processes and their Applications*, 87:167–197, 2000.

[Gopalsamy, 1992] K. Gopalsamy. *Stability and Oscillations in Delay Differential Equations of Population Dynamics*. Kluwer Academic Publisher, 1992.

[Gould and Lewontin, 1979] S.J. Gould and R. Lewontin. The spandrels of san marco and the panglossian paradigm: A critique of the adaptationist programme. *Proceedings of the Royal Society of London*, B205:581–598, 1979.

[Gyory and Ladas, 1991] I. Gyory and G. Ladas. *Oscillation Theory of Delay Differential Equations with Applications*. Oxford Science Publications, 1991.

[Han and Krogh, 2006] Zhi Han and B. H. Krogh. Reachability analysis of large–scale affine systems using low–dimensional polytopes. In J. Hespanha and A. Tiwari,

editors, *Hybrid Systems: Computation and Control*, volume 3927 of *Lecture Notes in Computer Science*, pages 287–301. Springer Verlag, 2006.

[Has'minskiy, 1980] R.Z. Has'minskiy. *Stochastic Stability of Differential Equations*. Sijthoff & Noordhoff, 1980.

[Hasty *et al.*, 1999] J. Hasty, J. Pradines, M. Dolnik, and J.J. Collins. Noise-based switches and amplifiers for gene expression. *PNAS*, 97(5):2075–2080, 1999.

[Hedlund and Rantzer, 2002] S. Hedlund and A. Rantzer. Convex dynamic programming for hybrid systems. *IEEE Transactions on Automatic Control*, AC-47(9):1526–1540, 2002.

[Henzinger *et al.*, 1998] T.A. Henzinger, P.W. Kopke, A. Puri, and P. Varaiya. What's decidable about hybrid automata? *Journal of Computer System Sciences*, 57:94–124, 1998.

[Hespanha *et al.*, 2001] J. Hespanha, S. Bohacek, K. Obraczka, and J. Lee. Hybrid modeling of TCP congestion control. In M.D. Di Benedetto and A. Sangiovanni-Vincentelli, editors, *Hybrid Systems: Computation and Control*, number 2034 in LNCS, pages 291–304. Springer-Verlag, 2001.

[Hespanha, 2004] J. P. Hespanha. Stochastic hybrid systems: Application to communication networks. In R. Alur and G. J. Pappas, editors, *Hybrid Systems: Computation and Control*, Lecture Notes in Computer Science 2993, pages 397–401. Springer Verlag, 2004.

[Hespanha, 2005] J. Hespanha. Polynomial stochastic hybrid systems. In M. Morari and L. Thiele, editors, *Hybrid Systems: Computation and Control*, LNCIS 3414, page 322338. Springer Verlag, 2005.

[Heymann *et al.*, 2002] M. Heymann, F. Lin, G. Meyer, and S. Resmerita. Analysis of Zeno behaviors in hybrid systems. In *Proceedings of the 41st IEEE Conference on Decision and Control*, Las Vagas, NV, 2002.

[Heymann *et al.*, 2005] M. Heymann, F. Lin, G. Meyer, and S. Resmerita. Analysis of Zeno behaviors in a class of hybrid systems. *IEEE Transactions on Automatic Control*, 50(3):376–383, 2005.

[Hu *et al.*, 2000] J. Hu, J. Lygeros, and S. Sastry. Towards a theory of stochastic hybrid systems. In N. Lynch and B. Krogh, editors, *Hybrid Systems: Computation and Control*, Lecture Notes in Computer Science 1790, pages 160–173. Springer Verlag, 2000.

[Hu *et al.*, 2003] J. Hu, M. Prandini, and S. Sastry. Probabilistic safety analysis in three-dimensional aircraft flight. In *Proceedings of the IEEE Conference on Decision and Control*, volume 5, pages 5335–5340, Maui, Hawaii, 2003.

[Hu *et al.*, 2004] J. Hu, W.-C. Wu, and S. Sastry. Modeling subtilin production in bacillus subtilis using stochastic hybrid systems. In R. Alur and G. J. Pappas, editors, *Hybrid Systems: Computation and Control*, Lecture Notes in Computer Science 2993, pages 417–431. Springer Verlag, 2004.

[Hu *et al.*, 2005] J. Hu, M. Prandini, and S. Sastry. Aircraft conflict prediction in the presence of a spatially correlated wind field. *IEEE Transactions on Intelligent Transportation Systems*, 3:326–340, 2005.

[Issacs, 1967] R. Issacs. *Differential Games*. John Wiley, 1967.

[Jacod and Shiryaev, 1987] J. Jacod and A.N. Shiryaev. *Limit Theorems for Stochastic Processes*. Springer Verlag, Berlin, 1987.

[Katoen, 2006] J.-P. Katoen. Stochastic model checking. In C.G. Cassandras and J. Lygeros, editors, *Stochastic hybrid systems*, Automation and Control Engineering Series 24, pages 79–106. Taylor & Francis Group/CRC Press, 2006.

[Kerrigan *et al.*, 2006] E. Kerrigan, D.Q. Mayne, and J Lygeros. Reachability analysis of discrete-time systems with disturbances. *IEEE Transactions on Automatic Control*, April 2006.

[Khalil, 2001] Hassan Khalil. *Nonlinear Systems (3rd edition)*. Prentice Hall, 2001.

[Kiesau *et al.*, 1997] P. Kiesau, U. Eikmanns, Z. Eckel, S. Weber, M. Hammelmann, and K. Entian. Evidence for a multimeric subtilin synthetase complex. *Journal of Bacteriology*, 179-5:1475–1481, 1997.

[Kitano, 2002] H. Kitano. Looking beyond the details: a rise in system-oriented approaches in genetics and molecular biology. *Current Genetics*, 41:1–10, 2002.

[Klapa, 2007] M. Klapa. Estimating the metabolic boundaries of zymomonas mobilis using linear programming analysis. In *Proceedings of the 2nd Conference on Foundations of Systems Biology in Engineering*, Stuttgart, DE, 2007.

[Klein and Entian, 1994] C. Klein and K Entian. Genes involved in self-protection against the lantibiotic subtilin produced by *bacillus subtilis* atcc 6633. *Applied and Environmental Microbiology*, 60(8):2793–2801, 1994.

[Kloeden and Platen, 1992] P.E. Kloeden and E. Platen. *Numerical Solution of Stochastic Differential Equations*. Springer Verlag, 1992.

[Kobayashi *et al.*, 2003] K. Kobayashi *et al.* Essential *Bacillus subtilis* genes. *PNAS*, 100(8):4678–4683, 2003.

[Kouretas *et al.*, 2006a] P. Kouretas, K. Koutroumpas, and J. Lygeros. Parameter identification for piecewise deterministic markov processes: a case study on a biochemical network. In *Proceedings of the 2nd IFAC Conference on Analysis and Design of Hybrid Systems*, Alghero, IT, 2006.

[Kouretas *et al.*, 2006b] P. Kouretas, K. Koutroumpas, J. Lygeros, and Z. Lygerou. Stochastic hybrid modelling of biochemical processes. In C. Cassandras and J.Lygeros, editors, *Stochastic Hybrid Systems – Automation and Control Engineering Series, 24*, volume 9083. CRC/Taylor & Francis, 2006.

[Kourzhanski and Varaiya, 1996] M. Kourzhanski and P. Varaiya. Stability of hybrid systems. In *Hybrid Systems III*, Lecture Notes in Computer Science 1066, pages 413–423. Springer Verlag, 1996.

[Koutsoukos and Riley, 2006] X. Koutsoukos and D. Riley. Computational methods for reachability analysis of stochastic hybrid systems. In J. Hespanha and A. Tiwari, editors, *Hybrid Systems: Computation and Control*, Lecture Notes in Computer Science 3927, pages 377–391. Springer Verlag, 2006.

[Krystul and Bagchi, 2004] Jaroslav Krystul and Arunabha Bagchi. Approximation of first passage time of switching diffusions. In *Proc. MTNS*, Leuwen, BG, July 2004.

[Kumar and Varaiya, 1986] P. R. Kumar and P. P. Varaiya. *Stochastic Systems: Estimation, Identification, and Adaptive Control*. Prentice Hall, Inc., New Jersey, 1986.

[Kurzhanski and Varaiya, 2000] A. B. Kurzhanski and P. Varaiya. Ellipsoidal techniques for reachability analysis. In N. Lynch and B.H. Krogh, editors, *Hybrid Systems: Computation and Control*, Lecture Notes in Computer Science 1790, page 202214. Springer Verlag, 2000.

[Kurzhanski and Varaiya, 2002] A. B. Kurzhanski and P. Varaiya. On reachability under uncertainty. *SIAM Journal of Control and Optimization*, 41(1):181–216, 2002.

[Kushner and Dupuis, 2001] H. J. Kushner and P.G. Dupuis. *Numerical Methods for Stochastic Control Problems in Continuous Time*. Springer-Verlag, New York, 2001.

[Kushner, 1984] H. J. Kushner. *Approximation and Weak Convergence Methods for Random Processes with Applications to Stochastic Systems Theory*. MIT Press, Cambridge, Massachussets, 1984.

[Kushner, 2002] H. J. Kushner. Numerical approximations for stochastic differential games. *SIAM Journal of Control and Optimization*, 41(2):457–486, 2002.

[Kveton *et al.*, 2006] B. Kveton, M. Hauskrecht, and C. Guestrin. Solving factored mdps with hybrid state and action variables. *Journal of Artificial Intelligence Research*, 27:1–49, 2006.

[Lafferriere *et al.*, 2000] G. Lafferriere, G. J. Pappas, and S. Sastry. O-minimal hybrid systems. *Mathematics of Control, Signals, and Systems*, 13(1):1–21, March 2000.

[Lafferriere *et al.*, 2001] G. Lafferriere, G. J. Pappas, and S. Yovine. Symbolic reachability computations for families of linear vector fields. *Journal of Symbolic Computation*, 32(3):231–253, September 2001.

[Lagunov, 1985] V.N. Lagunov. *Introduction to differential games and control theory*. Heldermann Verlag, Berlin, 1985.

[Liberzon *et al.*, 1999] D. Liberzon, J. Hespanha, and A.S. Morse. Stability of switched linear systems: a lie-algebraic condition. *Systems and Controls Letters*, 37(3):117–122, 1999.

[Liberzon, 2003] Daniel Liberzon. *Switching in Systems and Control*. Birkhauser, Boston, MA, 2003.

[Lincoln and Tiwari, 2004] P. Lincoln and A. Tiwari. Symbolic systems biology: hybrid modeling and analysis of biological networks. In R. Alur and G. J. Pappas, editors, *Hybrid Systems: Computation and Control*, Lecture Notes in Computer Science 2993, pages 660–672. Springer Verlag, 2004.

[Lygeros and Godbole, 1997] J. Lygeros and D.N. Godbole. An interface between continuous and discrete event controllers for vehicle automation. *IEEE Transactions on Vehicular Technology*, 46(1):229–241, 1997.

[Lygeros and Watkins, 2003] J Lygeros and O. Watkins. Stochastic reachability for discrete time systems: an application to aircraft collision avoidance. In *Proceedings of the 42nd Conference of Decision and Control*, Maui, Hawaii USA, 2003.

[Lygeros *et al.*, 1997] J. Lygeros, C.J. Tomlin, and S. Sastry. Multi-objective hybrid controller systhesis. In *Proceedings of the International Workshop on Hybrid and Real-Time Systems*, Lecture Notes in Computer Science 1201, pages 109–123. Springer Verlag, 1997.

[Lygeros *et al.*, 1998] J. Lygeros, D.N. Godbole, and S. Sastry. Verified hybrid controllers for automated vehicles. *IEEE Transactions on Automatic Control*, 43(4):522–539, 1998.

[Lygeros *et al.*, 1999] J. Lygeros, C. Tomlin, and S. Sastry. Controllers for reachability specifications for hybrid systems. *Automatica*, 35(3):349–370, 1999.

[Lygeros *et al.*, 2003] J. Lygeros, K. Johansson, S. Simic, J. Zhang, and S. Sastry. Dynamical properties of hybrid automata. volume 48 of *IEEE Transactions on Automatic Control*, 2003.

[Lygeros *et al.*, 2006] J. Lygeros, X. Mao, and C. Yuan. Stochastic hybrid delay population dynamics. In J. Hespanha and A. Tiwari, editors, *Hybrid Systems: Computation and Control*, Lecture Notes in Computer Science 3927, pages 436–450. Springer Verlag, 2006.

[Lygeros, 1996] J. Lygeros. *Hierarchical, Hybrid Control of Large Scale Systems*. PhD thesis, Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, 1996.

[Lygeros, 2004a] J. Lygeros. Lecture notes on hybrid systems. Technical report, ENSIETA, 2004.

[Lygeros, 2004b] J. Lygeros. On reachability and minimum cost optimal control. *Automatica*, 40 - 6:317–927, 2004.

[Lynch *et al.*, 2001] Nancy Lynch, Roberto Segala, and Frits Vaandrager. Hybrid i/o automata revisited. In M.D. Di Benedetto and A. Sangiovanni-Vincentelli, editors, *Hybrid Systems: Computation and Control*, LNCIS 2034, pages 403–417. Springer Verlag, 2001.

[Lynch *et al.*, 2003] Nancy Lynch, Roberto Segala, and Frits Vaandrager. Hybrid i/o automata. *Information and Computation*, 185(1):105–157, 2003.

[Malhame and Chong, 1985] R. Malhame and C-Y Chong. Electric load model synthesis by diffusion approximation of a high-order hybrid-state stochastic system. *IEEE Transactions on Automatic Control*, 30(9):854–860, 1985.

[Marcus, 1978] S. Marcus. Modeling and analysis of stochastic differential equations driven by point processes. *IEEE Transactions on Information Theory*, 24:164–172, 1978.

[Mariton, 1990] M. Mariton. *Jump Linear Systems in Automatic Control*. Marcel Dekker, New York, 1990.

[Maynard Smith, 1982] J. Maynard Smith. *Evolution and the Theory of Games*. Cambridge University Press, New York, 1982.

[McAdams and Arkin, 1997] H. McAdams and A. Arkin. Stochastic mechanisms in gene expresion. *PNAS*, 94:814–819, 1997.

[McAdams and Arkin, 1998] H. McAdams and A. Arkin. Simulation of prokaryotic genetic circuits. *Annual Reviews of Biophysical and Biomolecular Strictures*, 27:199–224, 1998.

[McAdams and Arkin, 1999] H. McAdams and A. Arkin. It's a noisy business! genetic regulation at the nanomolecular scale. *Trends in Genetics*, 15(2):65–69, 1999.

[Mikulevicius and Platen, 1988] R. Mikulevicius and E. Platen. Time discrete Taylor approximations for Ito processes with jump component. *Mathemathische Nachrichten*, 138:93–104, 1988.

[Milstein, 1994] G.N. Milstein. *Numerical Integration of Stochastic Differential Equations*. Kluwer Academic Publishing, 1994.

[Mitchell and Templeton, 2005] I. Mitchell and J. Templeton. A toolbox of hamilton-jacobi solvers for analysis of nondeterministic continuous and hybrid systems. In M. Morari and L. Thiele, editors, *Hybrid Systems: Computation and Control*, LNCIS 3414, pages 480–494. Springer Verlag, 2005.

[Mitchell and Tomlin, 2000] I. Mitchell and C. Tomlin. Level set methods for computation in hybrid systems. In B. Krogh and N. Lynch, editors, *Hybrid Systems: Computation and Control*, Lecture Notes in Computer Science, pages 310–323. Springer Verlag, 2000.

[Mitchell *et al.*, 2001] I. Mitchell, A. Bayen, and C. Tomlin. Validating a Hamilton-Jacobi approximation to hybrid system reachable sets. In M.D. Di Benedetto and A. Sangiovanni-Vincentelli, editors, *Hybrid Systems: Computation and Control*, Lecture Notes in Computer Science, pages 418–432. Springer Verlag, 2001.

[Mitchell *et al.*, 2005] I. Mitchell, A. M. Bayen, and C. Tomlin. A time-dependent hamilton-jacobi formulation of reachable sets for continuous dynamic games. *IEEE Transactions on Automatic Control*, 50(7):947–957, 2005.

[Mitchell, 2004] I. M. Mitchell. A toolbox for level set methods. Technical Report TR-2004-09, University of British Columbia, Vancouver, BC, Canada, 2004.

[Msadek, 1999] T. Msadek. When the going gets tough: survival strategies and environmental signaling networks in *bacillus subtilis*. *Trends in Microbiology*, 7-5:201–207, 1999.

[Nel, 1968] L.D. Nel. Theorems of stone-weierstrass type for non-compact spaces. *Mathematische Zeitschrict*, 104:226–230, 1968.

[Novak and Tyson, 1997] B. Novak and J.J. Tyson. Modeling the control of dna replication in fission yeast. *PNAS*, 94(17):9147–9152, 1997.

[Oishi *et al.*, 2001] M. Oishi, C. Tomlin, V. Gopal, and D. Godbole. Addressing multiobjective control: Safety and performance through constrained optimization. In M.D. Di Benedetto and A. Sangiovanni-Vincentelli, editors, *Hybrid Systems: Computation and Control*, LNCIS 2034, pages 459–472. Springer Verlag, 2001.

[Oksendal, 1998] Bernt K. Oksendal. *Stochastic Differential Equations: Theory and Applications - 5th Edition*. Sppringer Verlag, 1998.

[Picasso and Bicchi, 2005] B. Picasso and A. Bicchi. Control synthesis for practical stabilization of quantized linear systems. *Ren. Sem. Mat. Universita' di Torino*, 63(4):397–410, 2005.

[Pola *et al.*, 2003] G. Pola, M. L. Bujorianu, J. Lygeros, and M. D. Di Benedetto. Stochastic hybrid models: An overview. In *International Federation of Automatic Control Conference on Analysis and Design of Hybrid Systems*, Saint-Malo, France, 2003.

[Polderman and Willems, 1998] J. Polderman and J. Willems. *Introduction to Mathematical Systems Theory: A Behavioral Approach*. Springer Verlag, New York, 1998.

[Prajna *et al.*, 2004] S. Prajna, A. Jadbabaie, and G. J.Pappas. Stochastic safety verification using barrier certificates. In *Proceedings of the IEEE Conference on Decision and Control*, 2004.

[Prandini and Hu, 2006a] M. Prandini and J. Hu. A stochastic approximation method for reachability computations. In H.A.P. Blom and J. Lygeros, editors, *Stochastic hybrid systems: theory and safety applications*, Lecture Notes in Control and Informations Sciences 337, pages 107–139. Springer Verlag, 2006.

[Prandini and Hu, 2006b] M. Prandini and J. Hu. Stochastic reachability: Theory and numerical approximation. In C.G. Cassandras and J. Lygeros, editors, *Stochastic hybrid systems*, Automation and Control Engineering Series 24, pages 107–138. Taylor & Francis Group/CRC Press, 2006.

[Prandini *et al.*, 2000] M. Prandini, J. Hu, J. Lygeros, and S. Sastry. A probabilistic approach to aircraft conflict detection. *IEEE Transactions on Intelligent Transportation Systems*, 1(4):199–220, December 2000.

[Puri, 1995] A. Puri. *Theory of Hybrid and Discrete Event Systems*. PhD thesis, Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, 1995.

[Puterman, 1994] M.L. Puterman. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, Inc., 1994.

[Ramadge and Wonham, 1989] P.J.G. Ramadge and W.M. Wonham. The control of discrete event dynamical systems. *Proceedings of the IEEE*, 77(1):81–98, 1989.

[Rao and Arkin, 2001] C. Rao and A. Arkin. Control motifs for intracellular regulatory networks. *Annu. Rev. Biomed. Eng.*, pages 391–419, 2001.

[Rao *et al.*, 2002] C.V. Rao, D.M. Wolf, and A.P. Arkin. Control, exploitation and tolerance of intracellular noise. *Nature*, 420:231–237, 2002.

[Rosen, 1967] R. Rosen. *Optimality Principles in Biology*. Plenum Press, New York, 1967.

[Rubinow, 1975] S.I. Rubinow. *Introduction to mathematical biology*. John Wiley, 1975.

[Sastry, 1999] Shankar Sastry. *Nonlinear Systems, Analysis, Stability and Control*. Springer Verlag, 1999.

[Segre *et al.*, 2002] D. Segre, D. Vitkup, and G.M. Church. Analysis of optimality in natural and perturbed metabolic networks. *PNAS*, 99-33:15112–15117, 2002.

[Shampine and Thompson, 2000] L.F. Shampine and S. Thompson. Event location for ordinary differential equations. *Comput. Math. Appl.*, 39(5-6):43–54, 2000.

[Simic *et al.*, 2005] S. Simic, K. Johansson, J. Lygeros, and S. Sastry. Towards a geometric theory of hybrid systems. *Dynamics of Discrete, Continuous, and Impulsive Systems, Series B*, 12(5-6):649–687, 2005.

[Stein *et al.*, 2002] T. Stein, S. Borchert, P. Kiesau, S. Heinzmann, S. Kloss, C. Klein, M. Helfrich, and K Entian. Dual control of subtilin biosynthesis and immunity in *bacillus subtilis*. *Molecular Microbiology*, 44-2:403–416, 2002.

[Stein *et al.*, 2003] T. Stein, S. Heinzmann, P. Kiesau, B. Himmel, and K .Entian. The *spa*-box for transcriptional activation of subtilin biosynthesis and immunity in *Bacillus subtilis*. *Molecular Microbiology*, 47(6):1627–1836, 2003.

[Stein *et al.*, 2005] T. Stein, S. Heinzmann, S. Düsterhus, S. Borchert, and K .Entian. Expression and functional analysis of the subtilin immunity genes *spaIFEG* in the subtilin-sensitive host *Bacillus subtilis* mo1099. *Bacteriology*, 187(3):822–828, 2005.

[Stein, 2005] T. Stein. *Bacillus subtilis* antibiotics: structures, syntheses and specific functions. *Molecular Microbiology*, 56-4:845–857, 2005.

[Stone, 1948] M.H. Stone. The generalized weierstrass approximation theorem. *Mathematics Magazine*, 21(4):167–184, 1948.

[Strogatz, 2001] S.H. Strogatz. Exploring complex networks. *Nature*, 410:268–276, 2001.

[Stursberg and Krogh, 2003] O. Stursberg and B.H. Krogh. Effcient representation and computation of reachable sets for hybrid systems. In A. Pnueli O. Maler, editor, *Hybrid Systems: Computation and Control*, Lecture Notes in Computer Science 2623, pages 482–497. Springer Verlag, 2003.

[Tabuada *et al.*, 2004] P. Tabuada, G. Pappas, and P. Lima. Compositional abstractions of hybrid control systems. *Journal of Discrete Event Dynamical Systems*, 14(2):203–238, 2004.

[Tanner and Pappas, 2002] H. Tanner and G. Pappas. Simulation relations for discrete-time linear systems. In *Proceedings of the 15th International Federation on Automatic Control World Congress*, Barcelona, SP, 2002.

[Tiwari and Khanna, 2002] A. Tiwari and G. Khanna. Series of abstractions for hybrid automata. In C. J. Tomlin and M. R. Greenstreet, editors, *Hybrid Systems: Computation and Control*, volume 2289 of *LNCS*, pages 465–478. Springer Verlag, 2002.

[Tiwari and Khanna, 2004] A. Tiwari and G. Khanna. Nonlinear systems: Approximating reach sets. In R. Alur and G. J. Pappas, editors, *Hybrid Systems: Computation and Control*, Lecture Notes in Computer Science 2993. Springer Verlag, 2004.

[Tjalsma *et al.*, 2004] H. Tjalsma, E. Koetje, R. Kiewet, O. Kuipers, M. Kolman, J. van der Laan, R. Daskin, E. Ferrari, and S. Bron. Engineering of quorum-sensing systems for improved production of alkaline protease by *bacillus subtilis*. *Journal of Applied Microbiology*, 96:569–578, 2004.

[Tomlin *et al.*, 1998a] C. Tomlin, J. Lygeros, and S. Sastry. Synthesizing controllers for nonlinear hybrid systems. In T. Henzinger and S. Sastry, editors, *Hybrid Systems: Computation and Control*, Lecture Notes in Computer Science 1386, pages 360–373. Springer Verlag, 1998.

[Tomlin *et al.*, 1998b] C. Tomlin, G. J. Pappas, and S. Sastry. Conflict resolution for air traffic management : a study in multi-agent hybrid systems. *IEEE Transactions on Automatic Control*, 43(4):509–521, April 1998.

[Tomlin *et al.*, 2000] C. Tomlin, J. Lygeros, and S. Sastry. A game theoretic approach to controller design for hybrid systems. *Proceedings of the IEEE*, 88(7):949–970, July 2000.

[Tomlin *et al.*, 2003] C. Tomlin, I. Mitchell, A. Bayen, and M. Oishi. Computational techniques for the verification and controlof hybrid systems. *Proceedings of the IEEE*, 91(7):986–1001, 2003.

[Varaiya, 1975] Pravin Varaiya. The martingale theory of jump processes. *IEEE Transactions on Automatic Control*, 20(1):34–42, Feb 1975.

[Varaiya, 1993] Pravin Varaiya. Smart cars on smart roads: problems of control. *IEEE Transactions on Automatic Control*, 38(2):195–207, Feb 1993.

[Voigt *et al.*, 2005] C. Voigt, D. Wolf, and A.Arkin. The *Bacillus subtilis sin* operon: an evolvable network motif. *Genetics*, 169:1187–1202, 2005.

[Weibull, 1995] J. Weibull. *Evolutionary Game Theory*. The M.I.T. Press, Cambridge, MA, 1995.

[Witsenhausen, 1966] H. S. Witsenhausen. A class of hybrid-state continuous-time dynamic systems. *IEEE Transactions on Automatic Control*, 11(2):161–167, 1966.

[Wolf and Arkin, 2003] D.M. Wolf and A.P. Arkin. Motifs, modules and games in bacteria. *Current Opinions Microbiology*, 6:125–134, 2003.

[Wolf *et al.*, 2005] D. Wolf, V. Vazirani, and A.Arkin. Diversity in times of adversity: probabilistic strategies in microbial survival games. *J. theor. biol.*, 234(2):227–253, 2005.

[Xia, 1994] Aihua Xia. Weak convergence of markov processes with extended generators. *The Annals of Probability*, 22:2183–2202, 1994.

[Yazarel and Pappas, 2004] H. Yazarel and G. J. Pappas. Geometric programming relaxations for linear system reachability. In *Proceedings of the 2004 American Control Conference, Boston, MA*, June 2004.

[Yazarel *et al.*, 2004] H. Yazarel, S. Prajna, and G. J. Pappas. Sos for safety. In *Proceedings of the IEEE Conference on Decision and Control*, 2004.

[Yuan and Lygeros, 2005a] C. Yuan and J. Lygeros. On the exponential stability of switching diffusion processes. *IEEE Transactions on Automatic Control*, 50:1422–1426, 2005.

[Yuan and Lygeros, 2005b] C. Yuan and J. Lygeros. Stabilization of a class of stochastic differential equations with markovian switching. *Systems and Control Letters*, 54:819–833, 2005.

[Yuan and Lygeros, 2006] C. Yuan and J. Lygeros. Asymptotic stability and boundedness of delay switching diffusions. *IEEE Transactions on Automatic Control*, 51(1):171–175, 2006.

[Yuan and Mao, 2004] C. Yuan and X. Mao. Concergence of the Euler-Maruyama method for stochastic differential equations with markovian switching. *Mathematics and Computers in Simulation*, 64:223–235, 2004.

[Zhang *et al.*, 2001] J. Zhang, K. H. Johansson, J. Lygeros, and S. Sastry. Zeno hybrid systems. *International Journal of Robust and Nonlinear Control*, 11(5):435–451, 2001.

Index